

Exhibit W

Nos. 16-1742, -1743, -1744

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

NETLIST, INC.,

Appellant,

v.

DIABLO TECHNOLOGIES, INC.,

Appellee.

Appeal from the United States Patent and Trademark Office,
case nos. IPR2014-00882, -00883, -01011

CORRECTED BRIEF FOR APPELLANT NETLIST, INC.

MEHRAN ARJOMAND
MORRISON & FOERSTER LLP
707 Wilshire Boulevard
Los Angeles, California 90017
Telephone: (213) 892-5630

BRIAN R. MATSUI
SETH W. LLOYD
MORRISON & FOERSTER LLP
2000 Pennsylvania Avenue, N.W.
Washington, D.C. 20006
Telephone: (202) 887-8784

Counsel for Appellant Netlist, Inc.

CERTIFICATE OF INTEREST

Counsel for appellant Netlist, Inc. certifies the following:

1. The full name of every party or amicus represented by me is:

Netlist, Inc.

2. The name of the real party in interest (if the party named in the caption is not the real party in interest) represented by me is:

N/A

3. All parent corporations and any publicly held companies that own 10% or more of the stock of the party or amicus curiae represented by me are:

N/A

4. The names of all law firms and the partners or associates that appeared for the party or amicus now represented by me in the trial court or are expected to appear in this court are:

MORRISON & FOERSTER LLP: Mehran Arjomand, Erol C. Basol, David S. Kim, Seth W. Lloyd, Brian R. Matsui, Jean Nguyen, Jonathan Z. Statman

Dated: August 22, 2016

/s/ Brian R. Matsui

TABLE OF CONTENTS

CERTIFICATE OF INTEREST	i
TABLE OF AUTHORITIES	v
STATEMENT OF RELATED CASES	vii
JURISDICTIONAL STATEMENT	1
STATEMENT OF THE ISSUES.....	1
INTRODUCTION	2
STATEMENT OF THE CASE.....	4
A. The Challenged Patents	4
1. Background on memory modules	4
2. The inventions.....	8
3. The '150 and '536 patents' claims.....	15
B. The Prior Art	17
C. Proceedings Before the Board.....	21
1. The '882 IPR ('150 patent).....	22
2. The '1011 IPR ('150 patent)	25
3. The '883 IPR ('536 patent)	27
SUMMARY OF ARGUMENT	29
STANDARD OF REVIEW	32
ARGUMENT	33

I.	THE BOARD UNREASONABLY CONSTRUED “SELECTIVELY ELECTRICALLY COUPLING” TO REQUIRE SIMPLY SELECTING.....	33
A.	The Claims Require A Circuit Configurable To Electrically Couple Specified Data Lines In Response To A Selection.....	34
1.	The claim language supports only one meaning and the Board erred by adopting a different one	35
2.	The specification further shows that the Board’s construction unreasonably expands the claims to include configurations that the patent expressly distinguishes.....	37
3.	The Board violated established law by bypassing the intrinsic record and relying on extrinsic evidence to adopt a meaning that conflicts with intrinsic evidence.....	46
B.	Because The Board’s Unpatentability Conclusions For Claims 15-17 And 31-33 Of The ’150 Patent Turn On Its Erroneous Claim Construction, The Decisions Should Be Reversed.....	48
II.	THE BOARD ERRED IN CONCLUDING THAT THE REMAINING CHALLENGED CLAIMS OF THE ’150 AND ’536 PATENTS ARE UNPATENTABLE, BECAUSE IT BASED THOSE CONCLUSIONS ON ITS ERRONEOUS CLAIM INTERPRETATION	51
A.	The Board Erred By Importing Its Flawed Analysis Of “Selectively Electrically Coupling” Into Its Analysis Of The ’536 Patent.....	52
B.	The Board Also Erred By Sua Sponte Adopting And Then Applying A Construction Of “Selectively Isolating” In The ’150 Patent For The First Time In Its Final Written Decision.....	55
C.	The Board’s Conclusions Of Unpatentability For Obviousness Of All Claims Of The ’150 Patent Fall With Its Erroneous Claim Interpretation Because The Board Evaluated All Claims Identically	57

III. THE BOARD UNREASONABLY CONSTRUED “A CIRCUIT CONFIGURED TO BE MOUNTED ON A MEMORY MODULE” TO REACH PART OF A CIRCUIT CONFIGURED TO BE MOUNTED ON PART OF A MEMORY MODULE.....	59
CONCLUSION.....	63

TABLE OF AUTHORITIES

	Page(s)
CASES	
<i>Belden Inc. v. Berk-Tek LLC</i> , 805 F.3d 1064 (Fed. Cir. 2015)	56
<i>Cuozzo Speed Techs., LLC v. Lee</i> , 136 S. Ct. 2131 (2016)	32
<i>Dell Inc. v. Accelleron, LLC</i> , 818 F.3d 1293 (Fed. Cir. 2016)	56, 61
<i>Ethicon Endo-Surgery, Inc. v. U.S. Surgical Corp.</i> , 93 F.3d 1572 (Fed. Cir. 1996).....	36
<i>In re Gartside</i> , 203 F.3d 1305 (Fed. Cir. 2000).....	32
<i>In re Magnum Oil Tools Int'l, Ltd.</i> , --- F.3d ---, No. 15-1300, 2016 WL 3974202 (Fed. Cir. July 25, 2016).....	56
<i>In re Man Machine Interface Techs. LLC</i> 822 F.3d 1282 (Fed. Cir. 2016)	34, 45, 46
<i>In re Varma</i> , 816 F.3d 1352 (Fed. Cir. 2016).....	34
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005)	3, 34, 38, 47
<i>PPC Broadband, Inc. v. Corning Optical Commc'ns RF, LLC</i> , 815 F.3d 747 (Fed. Cir. 2016)	34
<i>SAS Inst., Inc. v. ComplementSoft, LLC</i> , --- F.3d ---, No. 15-1346, 2016 WL 3213103 (Fed. Cir. June 10, 2016)	56
<i>Teva Pharm. USA, Inc. v. Sandoz, Inc.</i> , 135 S. Ct. 831 (2015)	32, 47
<i>Teva Pharm. USA, Inc. v. Sandoz, Inc.</i> , 789 F.3d 1335 (Fed. Cir. 2015).....	47
<i>Thorner v. Sony Comput. Entm't Am. LLC</i> , 669 F.3d 1362 (Fed. Cir. 2012).....	38
<i>Trivascular, Inc. v. Samuels</i> , 812 F.3d 1056 (Fed. Cir. 2016)	38
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	34
<i>World Class Tech. Corp. v. Ormco Corp.</i> , 769 F.3d 1120 (Fed. Cir. 2014)	62

STATUTES

5 U.S.C. § 554(b)(3).....	56
28 U.S.C. § 1295(a)(4)(A)	1
35 U.S.C. § 102(e)	21
35 U.S.C. § 103.....	21, 24
35 U.S.C. § 311.....	1

STATEMENT OF RELATED CASES

These are appeals from final written decisions of the Patent Trial and Appeal Board (“Board”) in *inter partes* reviews (“IPR”) of claims 15-17, 22, 24, 26, and 31-33 of U.S. Patent No. 7,881,150 (“’150 patent”) and claims 1, 16, 17, 24, and 30-31 of U.S. Patent No. 8,081,536 (“’536 patent”). No appeal from these proceedings has previously been before this Court or any other court.

Netlist, Inc. has asserted the ’150 and ’536 patents in the following case:

Netlist, Inc. v. Smart Modular Technologies, U.S. District Court for the Northern District of California, Civil Action No. 4:13-CV-05889-YGR.

The ’150 patent also is at issue in the following Board proceeding: *Inter Partes* Review of U.S. Patent No. 7,881,150 filed April 7, 2015 (Case No. IPR2015-01020).

The ’536 patent also is at issue in the following Board proceeding: *Inter Partes* Review of U.S. Patent No. 8,081,536 filed April 7, 2015 (Case No. IPR2015-01021).

Counsel is aware of no other pending case that will affect or be affected directly by this Court’s decision.

JURISDICTIONAL STATEMENT

These are appeals from the final written decisions of the Board in three IPRs filed by petitioner-appellee Diablo Technologies, Inc. (“Diablo”). The Board had jurisdiction under 35 U.S.C. § 311. In all three appealed IPRs—IPR2014-00882, IPR2014-00883, and IPR2014-01011—the Board entered its final written decisions on December 14, 2015. Netlist filed timely notices of appeal on February 10, 2016. This Court has jurisdiction under 28 U.S.C. § 1295(a)(4)(A).

STATEMENT OF THE ISSUES

1. Whether the Board erred in construing a “circuit . . . configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line” by elevating technical dictionary definitions above the intrinsic record.
2. Whether the Board’s error in construing claims with the “selectively electrically coupling” term requires vacating and remanding the Board’s other unpatentability determinations because the Board imported its erroneous reasoning for that claim construction to find unpatentable claims that do not even recite that term.

3. Whether the Board erred in construing “a circuit configured to be mounted on a memory module” to encompass circuitry that is *not* mounted on a memory module.

INTRODUCTION

In these consolidated appeals from IPRs a single claim construction error tainted every unpatentability determination that the Board made. Relying on little more than dictionary definitions, the Board construed the term “selectively electrically coupling” to include making a “selection” between different hard-wired memory devices on a memory module. But the plain meaning of that phrase does not require simply a “*selection*”—it requires “coupling.” To selectively electrically *couple* data lines, as recited, is to electrically *connect* different data lines in a memory module selectively. This “connecting” is *selective* because it is not a permanent connection; sometimes the connection is made, sometimes not.

Had the Board considered the specification, it would have realized that its overbroad and unreasonable construction was wrong. The specification teaches that prior art conventional memory modules contained fixed connections. Rather than selectively electrically couple data lines, prior art modules could receive “chip select” signals, which designated which memory devices should receive or transmit data signals on the fixed connections at any given time. The specification further explained that these fixed connections created inefficiencies—they always expose

the computer system to the electrical load of the memory devices, even when data signals are not received or transmitted. Netlist overcame this deficiency in the prior art and implemented its solution as part of its inventions. When a switch or other mechanism is used to selectively electrically couple the data lines, the electrical load of the memory device can be disconnected from the computer system.

The Board's construction ignored these teachings. It permitted the claims to read on the precise type of prior art conventional memory modules that the specification expressly distinguished: those with permanently coupled data signal lines that constantly expose the computer system to a significant electrical load. The Board did so simply because the specification contained no express definition of "selectively electrically coupling." Yet claim construction does not limit consideration of the specification to where a patentee acts as his own lexicographer. This Court's precedent requires that the claims always be read in light of the specification. And ever since this Court's en banc decision in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005), it is error to resort first to a dictionary definition before considering the intrinsic record. But that is what the Board did here.

If that were not enough, the Board compounded its error when it imported its reasoning for this erroneous construction to other claims that do not even have the

“selectively electrically coupling” requirements. Independent claims of both patents require “selectively isolating/isolate” rather than “selectively electrically coupling.” But rather than independently analyze these claims or whether the prior art references disclose “selectively isolating,” the Board simply repeated its flawed reasoning regarding “selectively electrically coupling.” That reasoning is still wrong when applied to a completely different claim phrase.

The Board’s decisions should be vacated in their entirety and remanded.

STATEMENT OF THE CASE

A. The Challenged Patents

1. Background on memory modules

The ’150 and ’536 patents relate to a “memory module” of a computer system. APPX164 (col.1:30-33).¹ Computer systems may use random access memory (RAM) to provide short-term storage of smaller quantities of data for the operation of active software programs, such as an Internet web browser that needs relatively small amounts of short-term memory to open and view webpages.

As computer systems have become more advanced, there has been an increased demand for RAM capacity. For example, to boost performance of a

¹ The ’150 patent is the parent of the ’536 patent, and the two patents share the same specification. Specification citations are to the ’150 patent unless otherwise noted.

personal computer, a user may purchase additional RAM, discretely packaged as a “memory module” product:

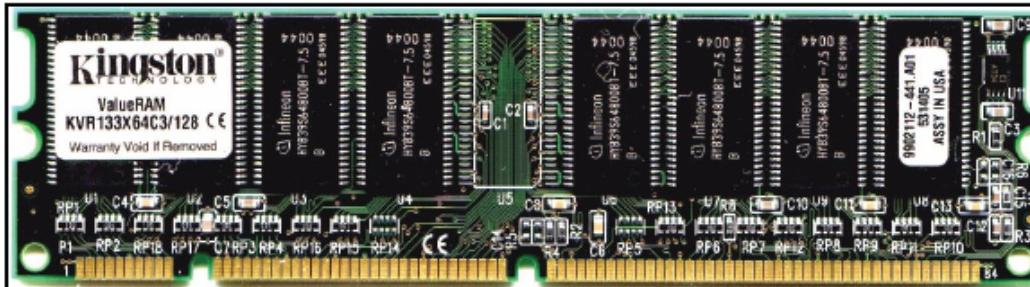


Figure 1: A Memory Module

A memory module is a computer board ("printed circuit board") with a handful of DRAM chips and associated circuitry attached to it. The picture is slightly larger than life-size.

APPX1917. The user may then open the housing of the personal computer and find a slot or socket in the motherboard dedicated to receive the memory module. APPX164 (col.1:35-40); *see also* APPX1237-1238 (¶25). By inserting additional RAM, or replacing existing RAM with new higher capacity RAM, the personal computer's performance will improve.

The computer system's motherboard has a CPU and a memory controller that are distinct components separate from any inserted memory modules. Hence, the CPU and memory controller are off the memory modules or, more simply, “off-module.” These distinct components are all shown in darker grey in the figure below:

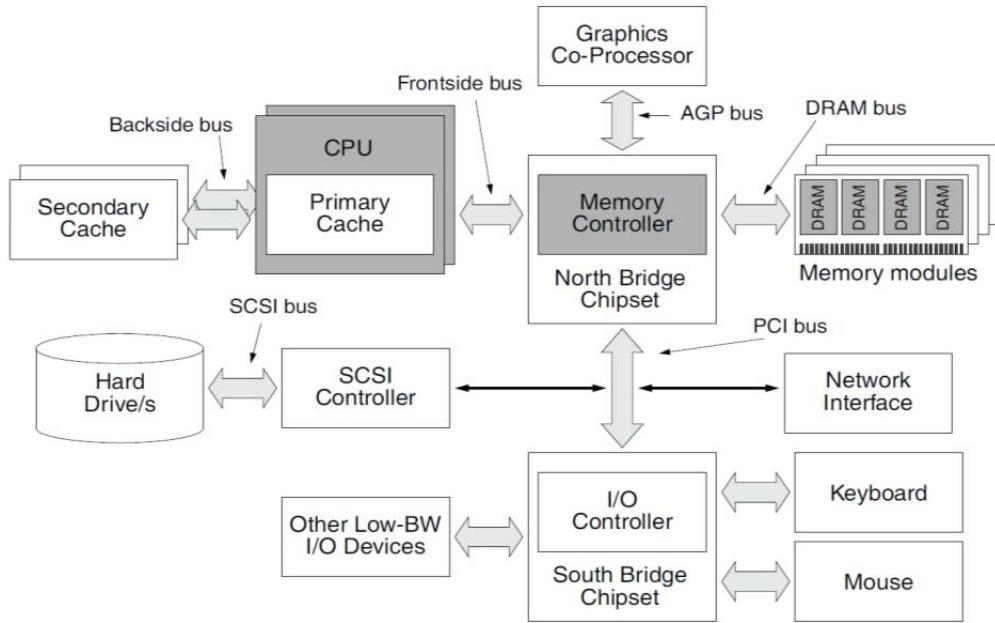


Figure 3: Typical PC Organization

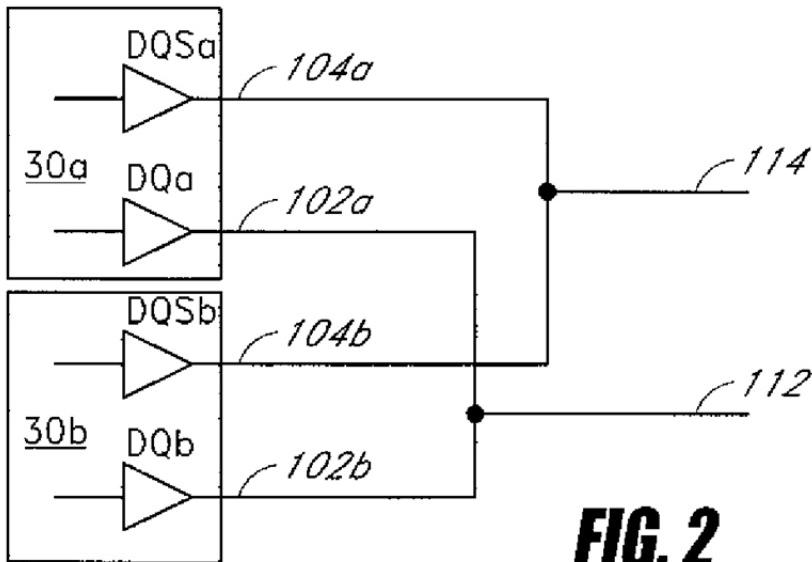
The DRAM subsystem is one part of a relatively complex whole. This figure illustrates a 2-way multiprocessor, with each processor having its own dedicated secondary cache. The parts most relevant to this report are shaded in darker grey: the CPU, the memory controller, and the individual DRAMs.

APPX1919. The memory controller acts as a liaison between the CPU and the RAM on the memory modules. *Id.*

The figures above show that each memory module contains multiple components called “memory devices” or “chips.” APPX1917, APPX1919. The figures show a specific type of memory device called “DRAM” or “DRAM chip.” “DRAM” stands for “dynamic” RAM. The claims at issue here are directed to a specific type of DRAM called Double Data Rate (DDR) Synchronous DRAM (SDRAM).

The memory devices on a memory module are often organized into groups called “ranks.” APPX164 (col.2:25-27). In a multi-rank memory module, multiple

ranks of DDR SDRAM devices share a common bus of data lines. This conventional arrangement is shown in FIG. 2 (below) of the '150 patent:



APPX142. On the left (**30a** and **30b**) are two memory devices on a memory module. APPX142, APPX166 (col.6:27-36). (In FIG. 2, each “rank” includes only a single memory device.) Each of the memory devices has a “DQ data signal line” associated with it (**102a** and **102b**). APPX142, APPX166 (col.6:27-36). “[E]ach of the memory devices **30a**, **30b** has their DQ data signal lines **102a**, **102b** electrically coupled to a common DQ line **112112** is in turn “electrically coupled” to the computer system’s memory controller (not shown). APPX166 (col.6:44-46).

This arrangement generates a load on the computer system. APPX168 (col.9:6-10:2). Electrical load refers to the physical effects that a component has when electrically connected to a circuit; such a load can be expressed in terms of

capacitance, inductance or impedance. APPX166 (col.5:3-5); APPX8720-8721 (¶68). An electrical load can affect circuit performance by drawing power or increasing signal delay. When memory devices' DQ data signal lines are permanently electrically connected to the common DQ signal line, as they are in the conventional module of FIG. 2, "the computer system is exposed to the loads of both memory devices **30a, 30b** concurrently." APPX166 (col.6:46-47).

2. *The inventions*

The claimed inventions of the '150 and '536 patents overcome several limitations of conventional memory modules to allow higher capacity modules and reduced costs. Two concepts are central to understanding how the patents achieve these advances: "memory density multiplication" and "load isolation/selective electrical coupling." *See, e.g.*, APPX174-180 (col.22:10-34:27) (section entitled "Memory Density Multiplication"); APPX165-169 (col.4:55-11:48) (section entitled "Load Isolation"). "Memory density multiplication" simulates a single memory device with a large number of memory locations (*e.g.*, a single 512-Mbit device) by using *multiple* memory devices with a lesser number of memory locations (*e.g.*, two 256-Mbit devices). "Load isolation/selective electrical coupling" ensures that the resulting increase in the number of memory devices on a module does not sacrifice performance.

The patents teach how to achieve both of these concepts without any change to the existing computer system, such as the memory controller.

Memory Density Multiplication On A Memory Module. To implement memory density multiplication, the memory module includes a circuit that translates between the “system memory domain” of the computer system and the “physical memory domain” of the memory devices or the memory module. Annotated FIG. 1 below of the ’150 patent illustrates a memory module **10** that performs translation:

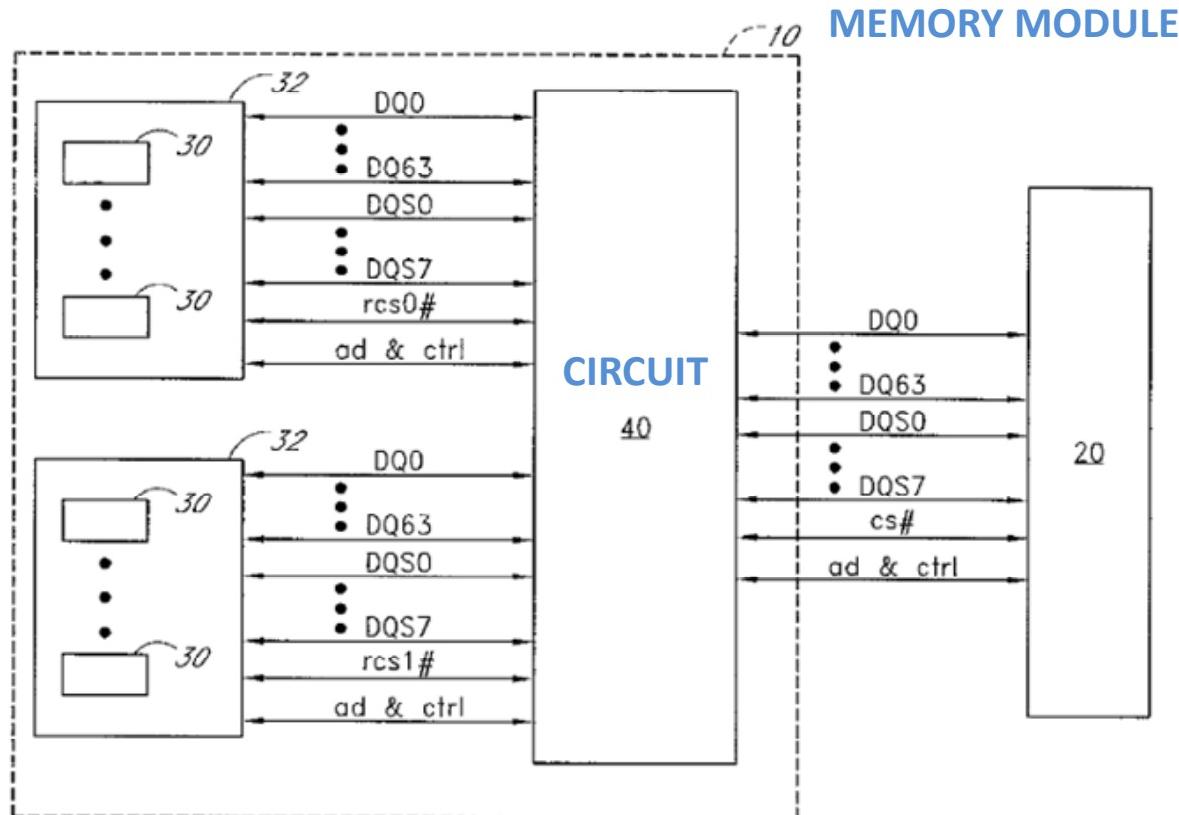


FIG. 1

APPX141. The memory module's circuit **40** receives input signals from the computer system's memory controller **20**. These input signals are based on a specific expected number of memory devices (e.g., a "virtual" memory module with eighteen 512-Mbit memory devices corresponding to a single rank). The circuit **40** translates these input signals into output signals that correspond to the actual number of physical devices on the memory module, which may be a larger number of devices than the computer system expects (e.g., thirty-six 256-Mbit memory devices corresponding to two ranks **32** of eighteen devices each).

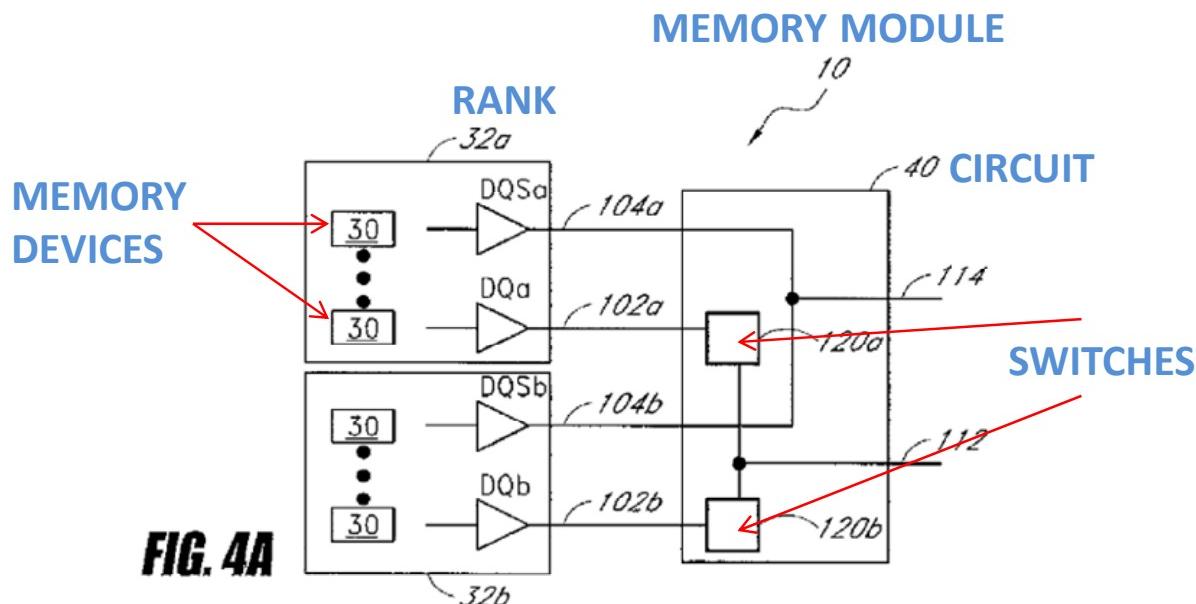
Memory density multiplication offers several advantages. It allows a single memory module to include a greater number of memory devices. The memory capacity of a memory module can be increased simply by adding more memory devices. APPX164 (col.2:32-42). Alternatively, by using multiple memory devices with fewer memory locations per device, memory density multiplication can provide a memory module with the same capacity but at a lower cost compared to other designs using fewer higher capacity memory devices. Memory devices are priced such that a larger number of such low-capacity memory devices can actually be cheaper than a smaller number of higher-capacity memory devices. APPX170 (col.14:45-67).

Load Isolation/Selective Electrical Coupling On A Memory Module.

Memory density multiplication is not without trade-offs, however. The "increased

load on the memory bus” due to the multiplied number of memory devices on the memory module “can degrade speed performance.” APPX168 (col.9:7-8). As shown above in FIG. 2 of the ’150 patent, each memory device electrically coupled to the common data line presents a load for the memory controller of the computer system. The ’150 patent addresses this issue by “selectively electrically coupling” data lines and selectively isolating the load(s) of individual or groups of memory devices on the memory module, thereby reducing the electrical load on the computer system. Such selective electrical coupling and isolating is implemented by the circuit used for memory density multiplication.

For example, annotated FIG. 4A below illustrates a memory module with a circuit for selectively electrically coupling a first data signal line to a common data signal line and for selectively electrically coupling a second data signal line to the common data signal line (or, conversely, selectively isolating the loads of memory devices that use the individual data signal lines):



APPX144 (item number **10** designates the memory module and refers to the entire schematic). As in FIG. 2, FIG. 4A shows multiple memory devices **30** on the left in two ranks **32a** and **32b**. On the right is the circuit **40** configured to be mounted on a memory module, such as for memory density multiplication from one rank in the system memory domain to two ranks **32a** and **32b** on the memory module **10**. APPX167 (col.7:59-65). Unlike FIG. 2, the circuit **40** of FIG. 4A includes two switches (**120a** or **120b**) for selectively connecting the device data signal lines (**102a**, **102b**) to the common data signal line (**112**). APPX167 (col.8:3-11). Switch **120a** can be “selectively actuated to selectively electrically couple” data signal line **102a** and common data signal line **112**. APPX167 (col.8:6-11). When switch **102a** is actuated, it electrically connects line **102a** and allows electric

current to flow to common line **112**. APPX167 (col.8:6-11). The same is true for switch **120b** and lines **102b** and **112**. APPX167 (col.8:6-11). When the switches are not actuated, they isolate the loads of the respective memory devices from the computer system. APPX167 (col.7:44-46). Thus, each switch allows selective electrical coupling between each respective pair of data lines and selective isolating of the loads of the corresponding memory devices. APPX167 (col.8:6-20, 48-53). Because the memory module's circuit **40** can selectively electrically couple each pair of data lines or selectively isolate the devices' loads, the total electrical load of the memory module that is seen by the computer system can be reduced. APPX166 (col.6:48-62), APPX167 (col.8:48-56).

Applying selective electrical coupling and isolating in the context of a memory module employing memory density multiplication was a challenge. In typical computer systems, the memory controller is aware of the number of memory devices and ranks of memory on a module. This allows the memory controller to avoid collisions or interference that might otherwise occur when, for example, successive read commands access multiple memory devices. APPX169 (col.12:60-67). Under memory density multiplication, however, the memory controller no longer knows the exact number of memory devices or ranks on a module, for example, that there are thirty-six memory devices on the module and not eighteen. As a result, collisions or interference can occur outside of the

memory controller's control. APPX170 (col.13:1-9). Avoiding these collisions thus becomes the responsibility of the memory density and load isolation circuit, which must coordinate timing events as part of its control over the switches that selectively connect memory devices to the memory controller.

Memory Density Multiplication and Load Isolation/Selective Electrical Coupling Combined On A Memory Module. Without selectively electrically coupling or some other form of load isolation, “the load of the memory module **10** may otherwise limit the number of ranks or the number of memory devices per memory module.” APPX167 (col.8:53-56). That is because “increased load on the memory bus can degrade speed performance” thereby presenting a roadblock to adding more memory devices for memory density multiplication. APPX168 (col.9:7-8). But by simultaneously increasing the ability to address a greater number of devices (memory density multiplication) and adding the capability to selectively electrically couple those devices to the computer system (load isolation and selective electrical coupling) “the speed *and* the memory density of the computer system are advantageously increased *without sacrificing one for the other.*” APPX168 (col.9:17-col.10:2) (emphases added).

3. *The '150 and '536 patents' claims*

The challenged claims of the '150 and '536 patents are directed to the combined concepts of “memory density multiplication” and “load isolation/selective electrical coupling” discussed above.

The '150 patent. Independent claim 15 of the '150 patent recites:

A circuit configured to be mounted on a memory module so as to be electrically coupled to a first double-data-rate (DDR) memory device having a first data signal line and a first data strobe line, to a second DDR memory device having a second data signal line and a second data strobe line, and to a common data signal line, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;

a register;

a phase-lock loop device configured to be operationally coupled to the first DDR memory device, the second DDR memory device, the logic element, and the register,

wherein the circuit is *configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line*, the circuit configurable to translate between the system memory domain of the computer system and a

physical memory domain of the memory module, wherein the system memory domain has a first memory density per memory device, and the physical memory domain has a second memory density per memory device less than the first memory density per memory device.

APPX184-185 (col.42:40-col.43:2) (emphases added highlighting disputed terms).

Independent claim 31 is similar. APPX185 (col.44:28-57). Whereas claims 15 and 31 require circuits that can “selectively electrically couple” data lines, independent claim 22 requires a circuit configurable to respond to input signals “by selectively isolating one or more loads of the DDR memory devices from the computer system.” APPX185 (col.43:51-54).

The '536 patent. Independent claim 1 recites:

A circuit configured to be mounted on a memory module configured to be operationally coupled to a computer system, the memory module having a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of DDR chip-select signals, the circuit including at least one configuration in which the circuit is configured to:

receive a set of signals comprising address signals and a second number of DDR chip-select signals smaller than the first number of DDR chip-select signals;

generate phase-locked clock signals and transmit the phase-locked clock signals to the DDR memory circuits of the first number of ranks;

selectively isolate a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals; and

generate the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of DDR chip-select signals.

APPX235 (col.41:20-43) (emphases added highlighting disputed terms).

Independent claim 24 claims a “method of operating a memory module” rather than a circuit. APPX236 (col.43:18-40). The method includes the step of “selectively isolating a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals.” APPX236 (col.43:33-36).

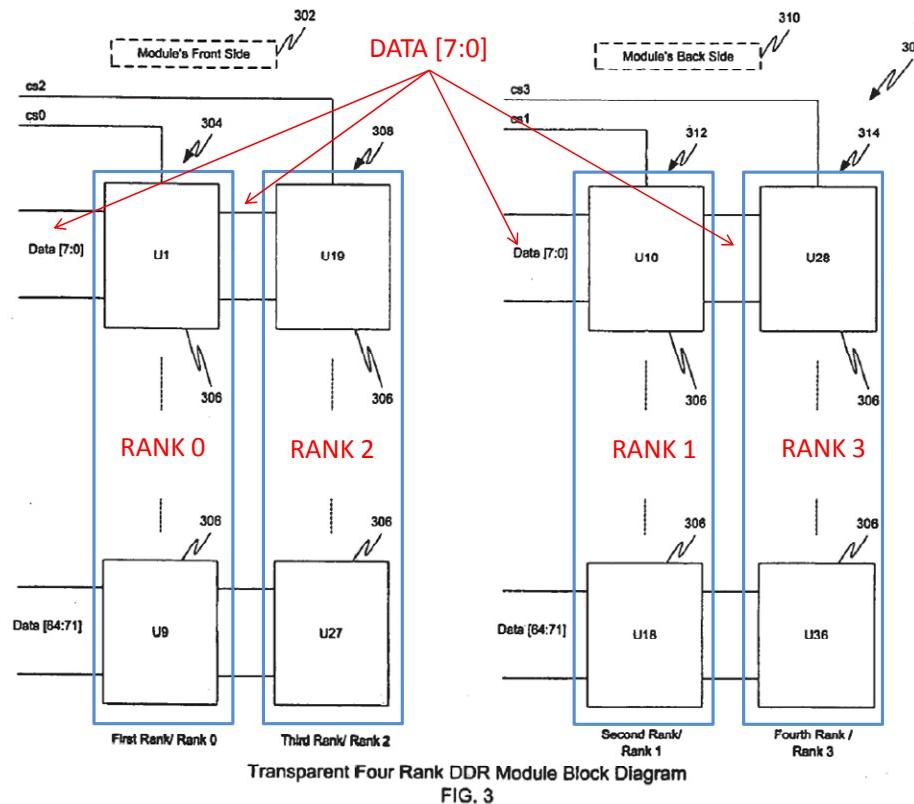
B. The Prior Art

The Board relied upon five references in finding the challenged claims of the ’150 and ’536 patents unpatentable.

Amidi. Amidi discloses “[a] transparent four rank memory module” for standard two-rank memory subsystems. APPX1611 (Abstract); APPX1622 (¶12).² As shown below in annotated FIG. 3, the memory module has four ranks of memory devices: ranks 0 and 2 are stacked on the front side of the memory

² The records from the three IPRs include many of the same documents. For simplicity and clarity, the record cites in this brief are to the first time a document appears in the appendix.

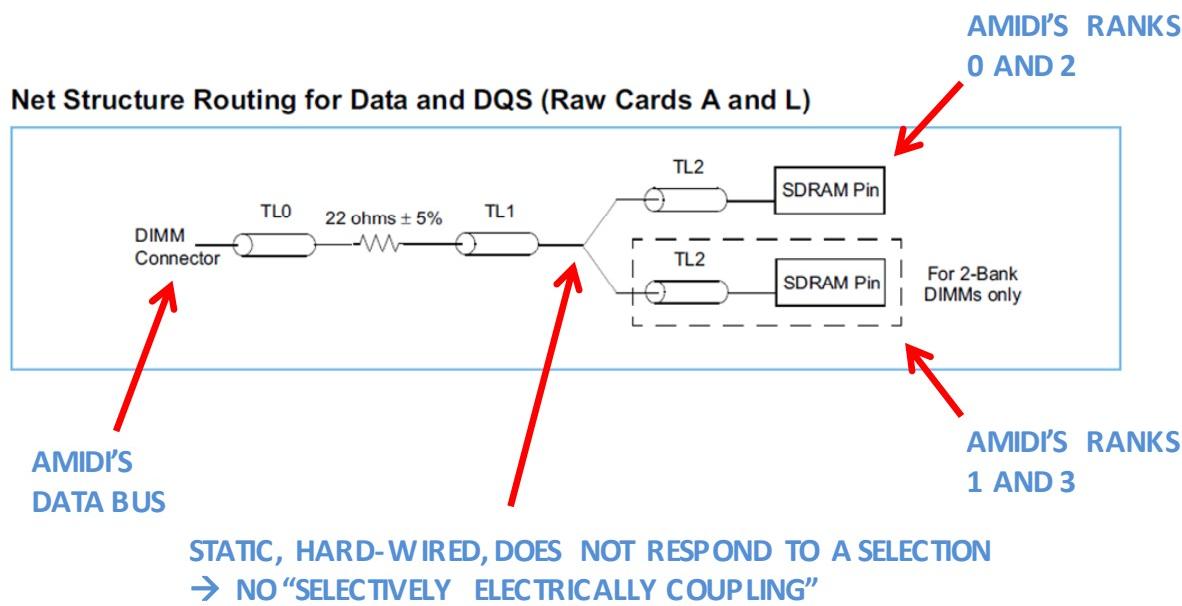
module, and ranks 1 and 3 are stacked on the back side of the memory module. The memory module includes a logic device (not shown in FIG. 3) that takes signals, such as chip-select signals, for a two-rank memory module and generates signals for the four-rank memory module. APPX1624 (¶41). In this way, Amidi purports to determine, or choose, one of the four ranks of memory devices on the memory module.



APPX1614 (annotated based on description at APPX1623-1624 (¶¶34-36)).

The memory module in FIG. 3 uses permanent, hard-wired data lines to connect the memory devices of the different ranks to a shared data bus. For example, data pins 0, 1, . . . , and 72 for each of the four ranks would be

respectively hard-wired to data bus lines 0, 1, . . . , and 72 of the 72-line data bus [71:0]. APPX1623 (¶23). That is, all the data lines of the four ranks are hard-wired to the same 72-line data bus. APPX1623-1624 (¶¶34-35). Amidi's module is thus wired similarly to JEDEC standard reference JEDEC21C 4-20-4, as illustrated in the annotated figure below:



APPX1852; *see also* APPX2501 (Netlist's expert explaining same); APPX2745-2748 (Diablo's expert confirming same). As shown, the Data and DQS pins of each of the SDRAM memory devices (or groups of memory devices) on the right are permanently connected by transmission lines (TL2, TL1, and TL0) directly to the “DIMM Connector,” which connects to the off-module circuitry.

Klein. Klein discloses reducing data bus capacitance by decoupling unaccessed memory modules from the data bus. APPX1628 (Abstract);

APPX1636 (¶23). The unaccessed memory modules are decoupled through switches, which are controlled by the memory controller. APPX1636 (¶28).

Wiggers. Wiggers is directed to “[a] memory system for minimizing the capacitive load of the memory data bus.” APPX1639 (Abstract). Like Klein, Wiggers discloses that the off-module memory controller operates switches that span across multiple memory modules. APPX1642-1643 (FIGs. 3 and 4); APPX1647 (col.4:61-65); APPX1648 (col.5:11-6:4). That is, the switches operate at a memory module-level, not at a device- or rank-level on a memory module. APPX2507 (¶88).

Ludwig. Ludwig discloses stacking memory chips and using interface circuitry (“VIC chip”) for buffering data. APPX4702 (col.3:29-33), APPX4702 (col.4:8-10). Such buffering may be provided by on-VIC chip buffers. APPX4704 (col.7:57-59). Importantly, Ludwig’s on-chip buffers buffer data transfers for the entire module, *i.e.*, module-level operation. APPX5602-5603 (¶83). Buffering may also be provided by off-VIC chip buffers. APPX4704 (col.7:33-37.) But off-VIC chip buffers are internal to the memory devices and not part of Ludwig’s VIC chip.

Dell. Dell is directed to a memory module that operates in a programmable addressing mode. APPX7689 (col.2:40-45). Dell provides for an ASIC that performs an address re-mapping operation. APPX7690 (col.4:24-32). Dell was

cited for its purported disclosure or suggestion of the recitations of dependent claims in the '536 patent.

C. Proceedings Before the Board

Diablo filed three separate petitions to institute *inter partes* review (IPR) of certain claims of Netlist's '150 and '536 patents. Two petitions (IPR2014-00882 and IPR2014-01011) were directed to the '150 patent and one (IPR2014-00883) to the '536 patent. The Board granted the petitions, instituting review on fifteen claims of the patents based on six grounds of unpatentability, as summarized in the table below:

IPR Number	Challenged claims	Asserted ground (and references)
IPR2014-00882	15-17, 22, 24, 26, and 31-33 of the '150 patent	§ 103 (Amidi and Klein)
IPR2014-00882	15-17, 22, 24, 26, and 31-33 of the '150 patent	§ 103 (Amidi and Wiggers)
IPR2014-01011	15-17, 22, 24, 26, and 31-33 of the '150 patent	§ 103 (Ludwig and Amidi)
IPR2014-01011	22, 24, and 26 of the '150 patent	§ 102(e) (Amidi)
IPR2014-00883	1, 16, 17, 24, and 30-31 of the '536 patent	§ 103 (Klein and Amidi)
IPR2014-00883	16, 17, and 30-31 of the '536 patent	§ 103 (Klein, Amidi, and Dell)

APPX2, APPX53, APPX99.

In three final written decisions, the Board concluded that all challenged claims of both patents were unpatentable. APPX50, APPX96, APPX134.

I. The '882 IPR ('150 patent)

The Board construed several claim terms in its final written decision.³ As relevant here, it adopted the following constructions:

Claim term	Construction	Affected claims
“A circuit configured to be mounted on a memory module”	“circuitry configured to be mounted on at least a portion of a memory module,” which encompasses “at least a portion of circuitry configured to be mounted on at least a portion of a memory module” APPX13-14	15-17, 22, 24, 26, and 31-33 of the '150 patent
“A circuit . . . wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line”	“making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component” APPX13	15-17 and 31-33 of the '150 patent

³ The Board’s final written decisions in all three IPRs largely tracked its institution decisions. Thus the summary here focuses on the final decisions, but notes any relevant differences from institution.

"A circuit . . . wherein the circuit is configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system"	No construction	22, 24, and 26 of the '150 patent
---	-----------------	-----------------------------------

Circuit Configured To Be Mounted On A Memory Module. Diablo did not offer a construction of "circuit configured to be mounted on a memory module." APPX13. It subsequently agreed with the Board's construction from the institution decision. APPX625-628. Netlist argued that the term means "an entire circuit configured to be mounted on a single memory module." APPX13.

The Board rejected Netlist's argument and construed the term as "circuitry configured to be mounted on at least a portion of a memory module." APPX13-14. The Board explained that the '150 patent "defines 'circuit' as 'a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.''" APPX14. According to the Board, that definition of "circuit" supported its construction. APPX14. The Board also clarified that its construction "encompass[es] 'at least a portion of circuitry configured to be mounted on at least a portion of a memory module.''" APPX14.

Selectively Electrically Coupling. Diablo again did not propose a construction for this term with its petition. APPX329-333. Netlist argued that the term meant “electrically coupling in response to a selection.” APPX11.

The Board adopted its own construction. It explained that the specification “does not define explicitly the term ‘selectively electrically coupling.’” APPX12. It pointed to the IEEE Dictionary, which “defines ‘electrical coupling’ as ‘[e]lectrical charges in conductors of a disturbed circuit formed by electrical induction.’” APPX12. Another part of the IEEE dictionary definition stated that “electrical coupling is also called capacitive coupling.” APPX12. The Board thus also looked to the IEEE Dictionary definition of “coupling capacitance (1) (ground systems)” which defined that term as “[t]he association of two or more circuits with one another by means of capacitance mutual to the circuits.” APPX12. Finally, it turned to the Oxford English Dictionary, which defines “selectively” as “[i]n a selective manner; by selection.” APPX12. Based on these three dictionary definitions, the Board adopted the construction “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component.” APPX13.

Obviousness. Based on its claim constructions, the Board concluded that the challenged claims were unpatentable under § 103 and in light of the asserted combinations of Amidi and Klein or Amidi and Wiggers. The Board emphasized

that its findings with regard to each of the disputed terms were based on its adopted constructions. For the “selectively electrically coupling” term, it pointed in part to Amidi’s disclosure of chip-select signals. The Board explained that, “as we have construed the term,” the claims reached “directing signals down a specific signal line or data bus in order to determine an active rank.” APPX29. That was so even if the lines on which signals were directed were “hard-wired.” APPX29. The Board also criticized Netlist for making arguments that the Board perceived to be based on Netlist’s proposed interpretations. APPX33.

2. *The ’101 IPR (’150 patent)*

The Board adopted the following relevant claim constructions:

Claim term	Construction	Affected claims
“A circuit configured to be mounted on a memory module”	“circuitry configured to be mounted on at least a portion of a memory module” which encompasses “a portion of circuitry configured to be mounted on at least a portion of a memory module” APPX64-66	15-17, 22, 24, 26, and 31-33 of the ’150 patent

<p>“A circuit . . . wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line”</p>	<p>“making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component” APPX64</p>	<p>15-17 and 31-33 of the ’150 patent</p>
<p>“A circuit . . . wherein the circuit is configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system”</p>	<p>“making a selection between at least two components and not transferring power or signal information from one selected component to the other selected component” (for ’1011 IPR)⁴ APPX64</p>	<p>22, 24, and 26 of the ’150 patent</p>

The Board’s final written decision for the ’1011 IPR largely tracked its decision in the ’882 IPR. It adopted the same constructions for “circuit configured to be mounted on a memory module” and the “selectively electrically coupling” term. APPX64. And it provided the same reasoning in support of these constructions. APPX64-66.

⁴ The Board stated that it was construing the term “selectively *electrically isolating*” from these claims. APPX64. But no claim of the ’150 patent (or the ’536 patent) contains that term.

The Board departed from its '882 decision, however, with regard to the "selectively isolating" term. In the '882 decision, the Board did not construe that term. Nor did it construe the term in its institution decision for the '1011 IPR. APPX3539-3545. Although neither party offered or argued for a construction, the Board's final written decision for the '1011 IPR construed the phrase to mean "making a selection between at least two components and not transferring power or signal information from one selected component to the other selected component." APPX64. According to the Board, that construction was proper "[b]ased on the same reasoning" it gave for its construction of "selectively electrically coupling." APPX64.

As in the '882 IPR, the Board relied on its interpretations of the disputed terms in reaching its unpatentability determinations.

3. *The '883 IPR ('536 patent)*

The Board's final written decision adopted the following relevant claim constructions:

Claim term	Construction	Affected claims
“A circuit configured to be mounted on a memory module”	“circuitry configured to be mounted on at least a portion of a memory module” which encompasses “at least a portion of circuitry configured to be mounted on at least a portion of a memory module” APPX106-108	1, 16, and 17 of the ’536 patent
<p>“A circuit . . . configured to . . . selectively isolate a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals”</p> <p>AND</p> <p>“A method of operating a memory module . . . comprising: . . . selectively isolating a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals”</p>	“electrical separation from one selected component from another selected component” APPX106	1, 16, 17, 24, and 30-31 of the ’536 patent

The Board adopted the same constructions for “circuit configured to be mounted on a memory module” that it had adopted in the final written decisions

for the '882 and '1011 IPRs. APPX106. And it again provided the same reasoning in support. APPX106-109.

For the “selectively isolate/isolating” terms of the '536 patent’s claims, the Board adopted the construction “electrical separation from one selected component from another selected component.” APPX106. That was the construction it had adopted during institution and which neither party challenged. APPX106.

In applying that construction to the combination of Klein and Amidi, the Board gave the same reasoning that it gave in the other IPRs. It again relied on Amidi’s disclosure of chip-select signals. APPX123-124. The Board explained that, “as we have construed the term,” the claims reached “directing signals down a specific signal line or data bus in order to determine an active rank.” APPX123. That was so even if the lines on which signals were directed were “hard-wired.” APPX123. The Board again criticized Netlist for making arguments that the Board perceived to be based on Netlist’s proposed interpretations. APPX127-128.

SUMMARY OF ARGUMENT

I. Relying on dictionary definitions and discounting the specification because it contained no express definition, the Board erred in interpreting a circuit “configurable to be responsive to the set of input signals by *selectively electrically coupling* the first data signal line to the common data signal line and *selectively electrically coupling* the second data signal line to the common data signal line” in

claims 15-17 and 31-33 of the '150 patent. The Board read the claims to encompass simply making a *selection* among memory devices with hard-wired, permanent connections. But that is not a reasonable construction. The claim language requires and the specification consistently describes that the circuit must electrically couple specified data lines in a selective manner. That requires a circuit that sometimes electrically connects the specified data lines and sometimes not.

The Board's rewriting of the claims is clear from how it explained its interpretation. The Board concluded that Amidi disclosed a memory module with the claimed feature because Amidi includes sending a signal to select an active memory device with "hard-wired, permanent data signal lines." But Amidi's disclosure is identical in this regard to the conventional memory module that the specification expressly distinguishes. The Board erred by reading the claims so broadly.

Because the Board's unpatentability conclusions with regard to claims 15-17 and 31-33 of the '150 patent turned on that construction, those determinations should be vacated.

II. The Board's error with regard to "selectively electrically coupling" tainted its analysis of the remaining challenged claims of both patents. Most of the other challenged claims require a circuit "configurable to be responsive to the set

of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system.” For both patents, the Board simply imported its flawed reasoning with regard to “selectively electrically coupling.” It read “selectively isolating” to reach simply making a selection so that a memory device was inactive on a hard-wired, permanent connection. But similar to “selectively electrically coupling,” selectively isolating cannot be achieved by merely selecting a device to be inactive. It requires electrically separating an electrical load in a selective manner, *i.e.*, sometimes electrically separating a load and sometimes not.

The Board’s analysis of the “selectively isolating” claims also fails because the Board denied Netlist’s procedural rights. It *sua sponte* adopted and applied a new construction of that term in its final written decision. Netlist never received any notice of the Board’s potential construction and had no opportunity to respond.

III. The Board also erred in construing “a circuit configured to be mounted on a memory module.” The Board interpreted that term to reach “circuitry configured to be mounted on at least a portion of a memory module.” It justified that interpretation because the word “circuit” has a broad meaning. But the Board was required to interpret the entire claim phrase, not just the word circuit. And the entire claim phrase requires a circuit “configured to be mounted on a memory module,” not a circuit configured to be mounted on “a portion” of a memory module.

The Board’s construction defeats an important purpose of the inventions. A key contribution of the inventions was to allow existing systems to be upgraded by simply swapping a conventional memory module for a module according to the invention. But the Board’s construction stretches the claims to cover circuits that cannot be implemented in that fashion—under the Board’s construction, upgrading a system would require re-wiring and other significant hardware alterations. The Board’s construction is unmoored from the claim language and the description of the inventions.

STANDARD OF REVIEW

This Court reviews the Board’s conclusions of law de novo and its findings of fact for substantial evidence. *In re Gartside*, 203 F.3d 1305, 1316 (Fed. Cir. 2000). Claim construction is a question of law that may involve underlying fact finding. *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 135 S. Ct. 831, 838 (2015). During an *inter partes* review, the Board gives claim terms their broadest reasonable interpretation consistent with the specification. *Cuozzo Speed Techs., LLC v. Lee*, 136 S. Ct. 2131, 2144-45 (2016).

ARGUMENT

I. THE BOARD UNREASONABLY CONSTRUED “SELECTIVELY ELECTRICALLY COUPLING” TO REQUIRE SIMPLY SELECTING

Claims 15-17 and 31-33 of the ’150 patent require, “[a] circuit . . . wherein the circuit is configurable to be responsive to the set of input signals by [1] selectively electrically coupling the first data signal line to the common data signal line and [2] selectively electrically coupling the second data signal line to the common data signal line.” APPX184-185 (col.42:40, 58-62; col.44:28, 46-50). By its express terms, the claims thus require electrically *coupling* two specified data lines in a selective manner. Yet the Board construed “selectively electrically coupling” to mean “making a *selection* between at least two components so as to transfer power or signal information from one selected component to at least the other selected component.” APPX13, APPX64 (emphasis added). In doing so, the Board erred by reducing “selectively electrically coupling” data lines to choosing, or “selecting,” between components, rather than electrically connecting them.⁵

⁵ For ease of discussion and unless otherwise noted, this brief refers to the claim phrase “A circuit . . . configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line” as the “selectively electrically coupling” limitation or phrase.

A. The Claims Require A Circuit Configurable To Electrically Couple Specified Data Lines In Response To A Selection

The Board’s construction violates several bedrock claim construction principles. First, the broadest-reasonable-interpretation standard requires that the Board’s construction be *reasonable*. If “there is only one reasonable meaning of the claim language, considered alone and in light of the specification,” it is unreasonable for the Board to adopt some other meaning. *In re Varma*, 816 F.3d 1352, 1359-60 (Fed. Cir. 2016); *see also PPC Broadband, Inc. v. Corning Optical Commc’ns RF, LLC*, 815 F.3d 747, 752 (Fed. Cir. 2016) (reversing Board construction that “fails to account for how the claims themselves and the specification inform the ordinarily skilled artisan as to precisely which ordinary definition the patentee was using”). Second, the Board errs when it adopts a construction so broad that it includes prior art examples that the specification expressly distinguishes. *In re Man Machine Interface Techs. LLC*, 822 F.3d 1282, 1286-87 (Fed. Cir. 2016). Third, extrinsic evidence—such as a dictionary definition—cannot be used to interpret a claim term contrary to the meaning shown by the intrinsic patent record. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1322-23 (Fed. Cir. 2005) (en banc) (citing *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1584-85 & n.6 (Fed. Cir. 1996)). The Board made all three errors here.

1. *The claim language supports only one meaning and the Board erred by adopting a different one*

The claim language itself points to only one reasonable meaning for a circuit “selectively electrically coupling” between two specified data lines: the circuit must be configurable to electrically *couple* the specified data lines in response to a selection. In response to the set of received input signals, therefore, the circuit must establish an electrical connection. And the claims identify *what* is to be connected—the first data signal line must be electrically connected to the common data signal line and the second data signal line must be electrically connected to the common data signal line. Finally, each electrical connection, or “coupling,” must be selective—the circuit must be capable of sometimes coupling the specified data lines and other times not.

The claims also make clear that selectively electrically coupling is different from permanent (or fixed) electrical connections. The claims refer to a permanent electrical connection using an entirely different phrase—“electrically *coupled*.” For example, the claims require a circuit mounted on a memory module “so as to be electrically coupled” to memory devices and a common data signal line. APPX184-185 (col.42:40-46; col.44:28-34). The memory module in turn must “be electrically coupled” to a memory controller. APPX184-185 (col.42:40-46; col.44:28-34). The claim language thus distinguishes “selectively electrically coupling” from the permanently hard-wired state of being “electrically coupled.”

See Ethicon Endo-Surgery, Inc. v. U.S. Surgical Corp., 93 F.3d 1572, 1579 (Fed. Cir. 1996) (holding that a particular claim's use of two different phrases dispositively showed two different meanings).

The Board's construction ignored this distinction. Instead, the Board adopted "making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component." APPX13, APPX64. But that construction focuses on "making a selection," such as selecting one memory device or another, including where there already may be fixed connections. It ignores that the required action in the phrase "selectively electrically coupling" is "coupling," not "selecting." Indeed, where the claims require "selecting" the patent uses the word "select." APPX184 (col.42:5-9) (circuit configurable to "select at least one rank"). The Board therefore got the construction wrong by focusing on the wrong word in the claim term.

Moreover, the Board's overbroad construction of "making a selection *between at least two components*," makes little sense given the claim language. Far from making a selection between any components, the claims specifically identify all of the relevant circuit parts that must be selectively electrically coupled—a first data signal line must be selectively electrically coupled to a common data signal line, or a second data signal line to the common data line. The

“selection” that the claims require is whether to electrically couple the identified data signal lines. The claim language leaves no room for selecting some unspecified components.

The final portion of the Board’s construction—“making a selection . . . so as to transfer power or signal information *from one selected component to at least the other selected component*” is also wrong. Indeed, that phrase is inconsistent with the first part of the Board’s construction. If a circuit must “make a selection between at least two components,” then the result would be only a single “selected” component. Yet the Board’s construction refers to two “selected” components—it requires transferring information “from one selected component to at least *the other* selected component.” Regardless, the very idea of selected components has no place in the claim, which explicitly identifies which two circuit elements must be coupled.

2. *The specification further shows that the Board’s construction unreasonably expands the claims to include configurations that the patent expressly distinguishes*

The specification further demonstrates the error in the Board’s construction. Yet the Board all but ignored the specification. In construing “selectively electrically coupling,” the Board’s only mention of the specification was to note that it “does not define explicitly the term.” APPX12, APPX63. That is not the proper analysis: “While the broadest reasonable interpretation standard is broad, it

does not give the Board an unfettered license to interpret the words in a claim without regard for the full claim language and the written description.” *Trivascular, Inc. v. Samuels*, 812 F.3d 1056, 1062 (Fed. Cir. 2016). Claims must always be read in the context of the specification and prosecution history. *Phillips*, 415 F.3d at 1315. And the intrinsic record is relevant beyond potentially providing an explicit re-definition, contrary to the Board’s implication (APPX12, APPX63). *E.g., Thorner v. Sony Comput. Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012) (claim interpretation always begins with the words of the claim “in the context of the specification and prosecution history”).

Here, the specification confirms that “selectively electrically coupling” carries its plain meaning in the context of the claim as a whole. The specification contrasts the operation of conventional memory modules that used hard-wired connections with the claimed selectively electrically coupling of data lines. But the Board’s construction reduced “selectively electrically coupling” to selecting among components with hard-wired connections. In doing so, the Board read the claims so broadly that they improperly reach the distinguished prior art conventional memory modules.

The ’150 patent explains that conventional memory modules did not include the claimed “selectively electrically coupling.” FIG. 2 illustrates a portion of a conventional memory module:

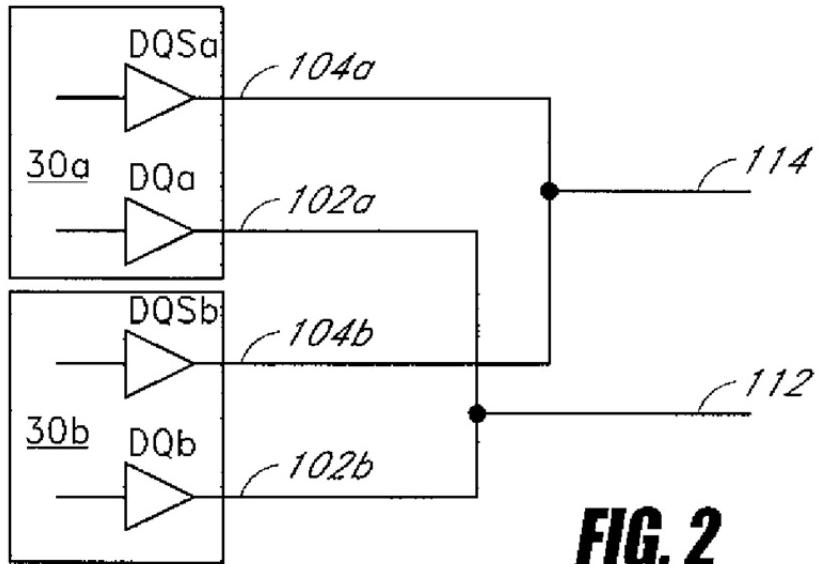
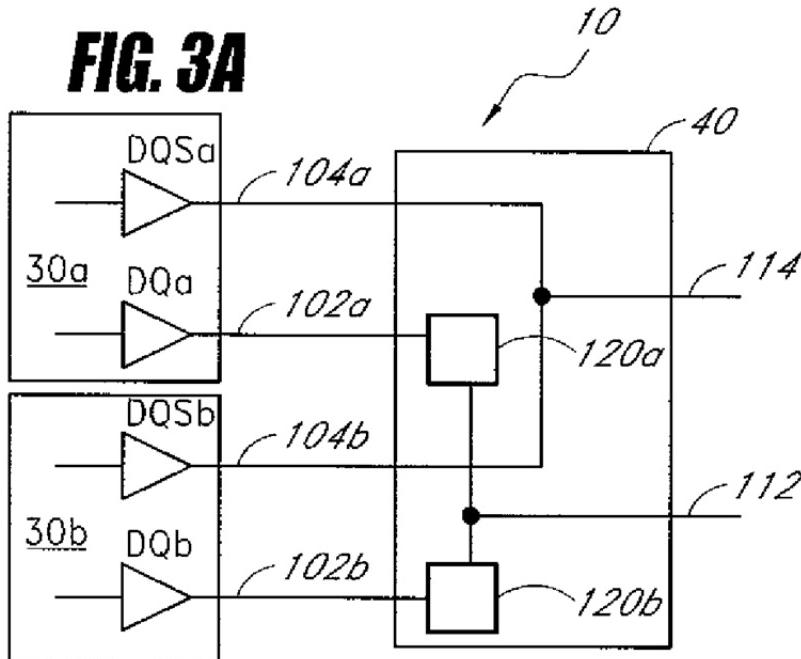


FIG. 2

APPX142. On the left (**30a** and **30b**) are two memory devices on a memory module. APPX142, APPX166 (col.6:27-36). Each device has a “DQ data signal line” associated with it (**102a** and **102b**). APPX142, APPX166 (col.6:27-36). And each device “has [its] DQ data signal line[] **102a**, **102b** electrically coupled to a common DQ line **112**.” APPX166 (col.6:39-43) (emphasis added). The common DQ, or data, line **112** is in turn “electrically coupled” to the computer system’s memory controller (not shown). APPX166 (col.6:44-46) (emphasis added). The memory controller controls which memory device is active on the common data line by sending a chip-select signal. APPX164 (col.2:43-47). But because both devices’ data lines remain permanently connected to the common data line, “the computer system is exposed to the loads of both memory devices **30a**, **30b** concurrently.” APPX166 (col.6:46-47).

In contrast, FIG. 3A illustrates a memory module according to the inventions with a circuit configurable for selectively electrically coupling two data lines:



APPX143 (item number **10** designates the memory module and refers to the entire schematic). Like FIG. 2, FIG. 3A shows two memory devices on the left. On the right is the circuit configured to be mounted on a memory module (item **40**). APPX167 (col.7:19-21). But unlike FIG. 2, the circuit **40** of FIG. 3A includes two switches (**120a** or **120b**) for selectively connecting the device data signal lines (**102a**, **102b**) to the common data signal line (**112**). APPX167 (col.7:19-26). Switch **120a** can be “selectively actuated to selectively electrically couple” data signal line **102a** and common data signal line **112**. APPX167 (col.7:22-26). When switch **120a** is actuated, it electrically connects line **102a** and allows electric current to flow to common line **112**. APPX167 (col.7:22-26). The same is true for

switch **120b** and lines **102b** and **112**. APPX167 (col.7:22-26). Thus, each switch allows selective electrical coupling between each respective pair of data lines. APPX167 (col.7:22-36); *see also* APPX143, APPX144, APPX145-146, APPX149-150, APPX166-167 (col.6:63-col.8:47), APPX170 (col.13:1-col.14:42) (illustrating and explaining other embodiments that operate in the same manner).

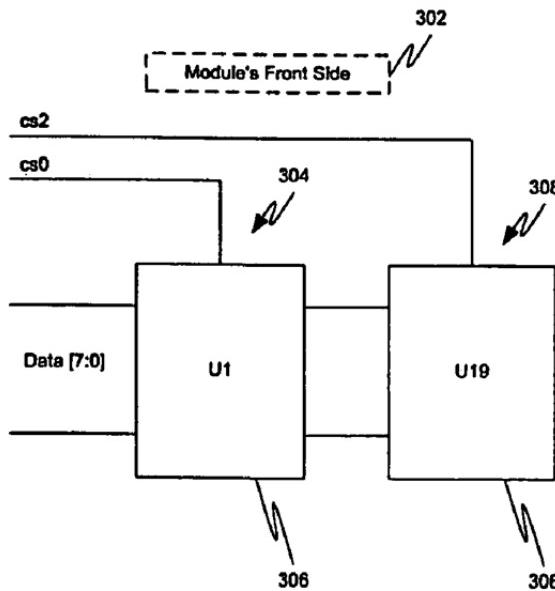
Adding the configurability to selectively electrically couple a device data line to a common data line was a key contribution of the inventions. When a memory device's data line is permanently electrically connected to the common data line, as is the case for both memory devices in the conventional memory module of FIG. 2, the computer system is always exposed to the electrical load of each such memory device on a module. APPX166 (col.6:46-47). But when a switch or other mechanism is used to selectively electrically couple the data lines, such as shown in FIG. 3A, the electrical load of the memory devices can be disconnected from the computer system by the switch. APPX166 (col.6:48-62), APPX167 (col.8:48-56). When the switch is not actuated, there is no electrical connection between the device data line and common data line, which means there is no electrical connection between the computer system (connected to the common data line) and the memory device (connected to the device data line). When that happens, the electrical load of the memory module is reduced to the

load of the circuit **40** (*i.e.*, the switches and other components of the circuit, but not the memory devices). APPX166 (col.6:48-63), APPX167 (col.8:48-56).

Solving the problem of increased loading due to additional memory devices was essential to overcoming the shortcomings of prior art memory modules. The capacity of prior art memory modules was limited by the number of memory devices that a single module could support. APPX164 (col.2:47-51). Two factors contributed to that limit: (1) memory controllers were designed to address a fixed maximum number of memory devices per module; and (2) increasing the number of memory devices on a module increased the electrical load on the system, which degraded speed performance and other system performance properties. APPX167-168 (col.8:53-col.9:9). To increase the number of memory devices that could be added onto a single module, the inventors simultaneously had to solve *both* of these problems. The inventors solved them by designing logic and circuit elements to perform memory density multiplication (thus allowing a computer system to address extra on-module memory devices) and to selectively electrically couple and load isolate individual or groups of memory devices (thus avoiding the increased electrical load of the extra memory devices).

The Board's construction ignores this context and obliterates the distinction between the conventional module in FIG. 2 and memory modules according to the inventions in FIG. 3A. This is evidenced from how the Board applied its

construction when concluding that the claims would have been obvious in light of Amidi. Netlist argued that Amidi did not disclose a circuit with the required “selectively electrically coupling” configurability. Rather, Amidi showed circuits with only permanent hard-wiring—electrically coupled, not “selectively electrically coupling”—as in a conventional module:



APPX1614 (FIG. 3 from Amidi, cropped to show only two memory devices). In particular, FIG. 3 from Amidi shows two memory devices, U1 and U19. APPX1614. Similar to the conventional memory module of FIG. 2 from the '150 patent, the data lines of Amidi's memory devices are directly and permanently coupled to a common data signal line: “As illustrated in **FIG. 3**, data bus [7:0] is connected to *both ranks' memory devices 306* (U1 and U19).” APPX1623 (¶34) (emphasis added); *see also* APPX2499-2500 (Netlist’s expert explaining same); APPX2743 (18:22) (Diablo’s expert testifying that “a common data bus, which is

the data 70” is “coupled to both U1 and U19”). Indeed, Amidi does not show any separation between a device data signal line and a common data line—it simply depicts a single data line bus (Data[7:0]) connected to both memory devices. And like conventional modules, Amidi discloses using a “chip select” signal for selecting which permanently coupled memory device should receive or transmit data signals via the common data line at any time. APPX1623 (¶34).

The Board nevertheless found that Amidi discloses the claimed circuit “configured to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line” based on its flawed construction. The Board explained that it was enough that Amidi “direct[s] signals down a specific signal line or data bus in order to determine an active rank within the memory devices.” APPX29, APPX83 (similar). But using chip-select signals when designating memory devices—as both Amidi and the conventional memory module in FIG. 2 do—is nothing more than *selecting* a device. Indeed, that was how the Board itself explained it—“Given our claim construction, we are unpersuaded by Patent Owner’s argument that the flow of electricity on *hard-wired, permanent data signal lines* does not constitute the electrical coupling of *selected components*.” APPX83 (emphases added); *see also* APPX85 (crediting Diablo’s expert’s testimony that Ludwig

discloses “‘multiple acts of coupling’” by “‘selecting the active chip’” (quoting APPX5369 (¶55))).

The Board’s reading of the claims is unreasonably broad. Directing signals down a specific line (*e.g.*, sending chip-select signals) to select a memory device and using hard-wired, permanent data lines is what the specification explains was done in a conventional memory module. The claims and written description use the word “transmit” to refer to directing signals and describe such hard-wired connections as “electrically coupled.” APPX166-167 (col.6:39-46, col.6:63-7:18). They distinguish these permanent connections from circuits configured for “selectively electrically coupling” pairs of data lines. Only the latter circuit feature furthers one of the key purposes of the inventions—reducing the electrical load on the computer system. APPX166 (col.6:46-62).

The Board’s construction is thus like the construction that this Court rejected in *Man Machine Interface*. There, the patent claimed a remote control device with “*a body adapted to be held by the human hand.*” 822 F.3d at 1284. The Board construed that phrase to reach the body of a desk-bound computer mouse. *Id.* at 1285. This Court held that the Board’s construction was unreasonably broad. The patent’s written description highlighted that the claimed device is “conveniently held in the user’s left or right hand,” making it “thus free to be held and moved around.” *Id.* at 1286 (quoting the disputed patent). And the description touted the

benefits of the device, including its “ease of use as the device is not desk bound, while being used or not used, [and] therefore can be kept and held in more user convenient positions and postures.” *Id.* (alteration in *Man Machine Interface*; again quoting the patent). The specification thus “expressly distinguishe[d] the remote control device from a desk-bound device” like a standard computer mouse. *Id.* It was error for the Board to adopt a construction that included this distinguished feature. *Id.*

The Board committed the same error here. Under its construction, the Board read Amidi as disclosing the claimed “selectively electrically coupling” of data signal lines. But Amidi discloses only using chip-select signals in selecting which permanently connected data device is active on a shared data line. And that disclosure is the same as the conventional memory module that the ’150 patent expressly distinguishes.

3. *The Board violated established law by bypassing the intrinsic record and relying on extrinsic evidence to adopt a meaning that conflicts with intrinsic evidence*

Eschewing the claims and the specification, the Board instead relied upon extrinsic evidence when construing “selectively electrically coupling.” The Board arrived at its construction based on a dictionary definition that equated “electrical coupling”—not “electrically coupling”—with “capacitive coupling.” APPX12, APPX63. It then used the dictionary definition of a term that does not appear

anywhere in the patents—“coupling capacitance (1) (ground systems)”—to reach its conclusion. APPX12 (citing APPX3264).

But that violates established precedent. *Phillips*, 415 F.3d at 1322-23. The Board’s reliance on a dictionary definition for “capacitive coupling” is particularly inappropriate because the specification never mentions or shows coupling “‘by means of capacitance mutual to the circuits,’” the definition the Board thought significant. APPX12, APPX63. Instead, all of the examples use conductive (*i.e.*, resistive) elements for selective electrical coupling, such as “field-effect transistor (FET) switches.” APPX166 (col.6:19-26), APPX167 (col.8:34-38), APPX168 (col.10:14-16), APPX170 (col.14:20-24). The Board’s use of an irrelevant dictionary definition provides no support for its construction and only confirms that it misunderstood the claims.⁶

⁶ The Board’s use of dictionary definitions does not trigger the deferential standard of review afforded subsidiary factual findings upon which claim construction sometimes depends, *see Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 135 S. Ct. 831, 841-42 (2015). The only “finding” the Board made was that the IEEE Dictionary and Oxford English Dictionary included certain definitions. APPX12, APPX63. There is no dispute about that. Determining whether a person of ordinary skill in the art would ascribe any of those meaning to the claim language in the context of these patents, however, is a question of law. *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 789 F.3d 1335, 1342 (Fed. Cir. 2015). This Court resolves that issue *de novo*. *Id.*

B. Because The Board’s Unpatentability Conclusions For Claims 15-17 And 31-33 Of The ’150 Patent Turn On Its Erroneous Claim Construction, The Decisions Should Be Reversed

All of the Board’s unpatentability conclusions for claims 15-17 and 31-33 of the ’150 patent fall with its erroneous claim construction. The Board’s decisions leaves no doubt that it held these claims unpatentable because it construed “selectively electrically coupling” of data lines to reach “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component.” Because that construction is wrong, the Board’s determinations of obviousness are also wrong.

The Board’s obviousness determinations were based on three different combinations of art: Amidi and Klein, Amidi and Wiggers, and Ludwig and Amidi. APPX50, APPX96. For each of these grounds, Diablo (and the Board) identified Amidi as disclosing the claim requirement of a circuit “configurable to be responsive to the set of input signals by selectively electrically coupling” specified data lines. APPX29, APPX47, APPX83. And for each of these grounds, the Board concluded that Amidi meets this requirement because it discloses “the flow of electricity on hard-wired, permanent data signal lines.” APPX83 (Ludwig and Amidi); *see* APPX29 (similar for Amidi and Klein), APPX47 (Board explaining that “both [Amidi and Wiggers] describe coupling or isolating memory device loads”). The Board further explained that Amidi meets the “*selectively*

“electrically coupling” requirement because it discloses using chip-select signals in selecting which device is active on those hard-wired, permanent data signal lines. APPX29, APPX47, APPX83. Based on its construction of “selectively electrically coupling,” the Board concluded that those disclosures were enough. APPX29 (Board relying on its construction to support that conclusion), APPX83 (same); *see* APPX47 (Board repeating that Amidi “describe[s] coupling or isolating memory device loads”).

The Board also looked at the disclosures of the other references in each of the combinations. APPX28-29 (discussing Klein), APPX48 (discussing Wiggers), APPX83 (discussing Ludwig). But to the extent that the Board examined whether those references disclosed the “selectively electrically coupling” limitation, its analysis again turned on its erroneous construction. For the Amidi-Klein combination, the Board primarily relied on Amidi to disclose the claim requirement. APPX29. The Board did note—after again emphasizing that the analysis was “under its construction” and in the context of a combination with Amidi—that Netlist’s expert testified that “Klein discloses the use of MOSFET switches.” APPX29. Yet the Board failed to elaborate or otherwise find that Klein disclosed the “selectively electrically coupling” limitation. APPX29. And it certainly never did so under Netlist’s proposed construction, nor with a proper view of how the correctly construed term fits with the other claim requirements,

see infra Section III (discussing requirement that “selectively electrically coupling” be performed by a “circuit configured to be mounted on a memory module”). APPX29.

With regard to the combination of Amidi and Wiggers, the Board simply recited the claim language and stated that the combination would have included the claim limitation, because “both references describe coupling or isolating memory device loads.” APPX47. Nowhere did it perform any analysis independent from its erroneous claim interpretation or under the proper construction of “selectively electrically coupling” in the context of the claim as a whole, *see infra* Section III. APPX47-48. Instead, the Board discounted Netlist’s arguments because it read them as “based on the claim constructions proffered by the Patent Owner and not on the constructions set forth in the Decision to Institute.” APPX48.

Finally, the Board similarly analyzed the combination of Ludwig and Amidi. APPX82-83. The Board found that Ludwig also disclosed the “selectively electrically coupling” requirement. APPX83. But as it did with Amidi, it did so based on its erroneous construction. Specifically, the Board “credit[ed] the testimony of Dr. Jagannathan, who states that Ludwig discloses VIC chip-select signals and address signals as ‘selectively electrically coupling.’” APPX83. Based on its construction, the Board thus thought it enough that Ludwig discloses “chip-select signals that selectively enable one of four memory layers and

selectively isolate the other memory devices in other ranks.” APPX83. But for the reasons given above, the claims require more than using chip-select signals to select an active memory device. *Supra*, Section I.A.

The Board’s unpatentability determinations with respect to claims 15-17 and 31-33 of the ’150 patent thus should be vacated and remanded for the Board to evaluate the patentability of these claims under the proper construction.

II. THE BOARD ERRED IN CONCLUDING THAT THE REMAINING CHALLENGED CLAIMS OF THE ’150 AND ’536 PATENTS ARE UNPATENTABLE, BECAUSE IT BASED THOSE CONCLUSIONS ON ITS ERRONEOUS CLAIM INTERPRETATION

The Board’s error in interpreting “selectively electrically coupling” infected its analysis of claims involving two other terms. All of its unpatentability determinations should accordingly be vacated.

All of the remaining challenged claims of the ’150 and ’536 patents require either a circuit “configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system” or a circuit configured to “selectively isolate a load of the DDR memory circuits of at least one rank . . . from the computer system in response at least in part to the set of signals.” APPX185 (43:51-54) (claim 22 of the ’150 patent), APPX235 (col.41:36-39) (claim 1 of the ’536 patent), APPX236

(col.43:33-36) (claim 24 of the '536 patent).⁷ “Selectively isolating/isolate” requires the complement of “selectively electrically coupling,” when that term is correctly construed. Instead of “coupling” the specified component, the component or its load must be isolated. Because the Board incorrectly construed “selectively electrically coupling,” it also incorrectly construed and misapplied “selectively isolating.” Thus, because the Board’s decisions for all of the challenged claims turned, at bottom, on its erroneous interpretation of “selectively electrically coupling,” all of its findings of unpatentability should be vacated.

A. The Board Erred By Importing Its Flawed Analysis Of “Selectively Electrically Coupling” Into Its Analysis Of The '536 Patent

The Board anchored its unpatentability conclusions for the '536 patent to its flawed reading of “selectively electrically coupling.” It construed “selectively isolate”/“selectively isolating” in the '536 patent to mean “electrical separation from one selected component from another selected component.” That construction and the Board’s understanding of it are wrong, as Netlist consistently argued. APPX6526-6527, APPX6683-6687.

⁷ Claim 24 of the '536 patent claims a method of operating a memory circuit that includes the step of “selectively isolating.” The Board failed to analyze that claim and its dependent claims separately from the circuit claims. APPX104 (using claim 1 as “illustrative); APPX129.

The Board's final written decision on the '536 patent demonstrates that it interpreted "selectively isolate" based on its interpretation of "selectively electrically coupling" in the '150 patent. When applying its interpretation of the '536 patent to the combination of Klein and Amidi, the Board simply imported its flawed reasoning from its final written decisions on the '150 patent with respect to "selectively electrically coupling."

For example, with regard to Amidi disclosing the "selectively isolate" element of the claims of the '536 patent, the Board said:

We do not agree with Patent Owner's position that hard-wired data signal lines cannot be electrically isolated in a selective fashion. . . . Amidi's disclosure of directing electrical signals down a specific signal line or data bus in order to electrically activate a rank within the memory devices and electrically inactivate other ranks falls within the scope of the term 'selectively isolate' as we have construed the term.

APPX123 (omission simply excludes the Board's recitation of its construction). That essentially replicates its analysis of "selectively electrically coupling" in the '150 patent. With regard to whether combinations with Amidi disclosed "selectively electrically coupling," the Board said:

Additionally, we are unpersuaded by Patent Owner's argument that hard-wired data signal lines, such as that taught by Amidi, cannot be electrically coupled in a selective fashion. . . . Amidi's disclosure of directing signals down a specific signal line or data bus in order to determine an activate a rank within the memory devices

falls within the scope of the term ‘selectively electrically coupling’ as we have construed the term.

APPX29 (same omission).

The Board’s construction of “selectively isolating/isolate” is wrong because the Board was wrong about what “selectively electrically coupling” means. As explained above, “selectively electrically coupling” does not require simply making a selection. *Supra*, Section I. It requires making electrical connections in a selective manner, *i.e.*, sometimes establishing an electrical connection and sometimes not. The claims and specification use the phrase “selectively isolating” similarly. The specification expressly states that “isolation” includes “electrical separation of one or more components from another component or from one another.” APPX166 (col.5:5-8). “Selectively isolating” therefore requires electrically separating one or more electrical loads of the memory devices from the computer system in a selective manner, *i.e.*, sometimes isolating the loads and sometimes not. But the Board again reduced “selectively isolating” to simply selecting—selecting a memory device or rank of memory devices to be inactive by using chip-select signals. The Board thus erred by once again incorrectly reading the claims to require simply “selecting.”

B. The Board Also Erred By Sua Sponte Adopting And Then Applying A Construction Of “Selectively Isolating” In The ’150 Patent For The First Time In Its Final Written Decision

The Board also erred in its anticipation analysis of the “selectively isolating” claims of the ’150 patent. It denied Netlist’s procedural rights by adopting a claim construction for “selectively isolating” in the ’150 patent for the first time in its final decision. Regardless, the Board’s claim interpretation is wrong for the same reason it was wrong with regard to the nearly identical term in the ’536 patent.

The Board held that claims 22 and 26 of the ’150 patent were unpatentable because they were anticipated by Amidi. APPX96. These claims require a circuit “configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system.” APPX185 (col.43:51-54). In concluding that these claims would have been anticipated, the Board *sua sponte* adopted a construction for that limitation *for the first time* in its final decision. Compare APPX62-64 (final decision adopting a construction for “selectively electrically isolating” based “on the same reasoning” as its construction for “selectively electrically coupling”), *with* APPX3539-3545, APPX3559-3560 (institution decision adopting no construction and only mentioning the limitation once, in reference to what Netlist argued).⁸ Neither party

⁸ The Board construed “selectively *electrically* isolating,” a phrase never used in the ’150 patent.

ever proposed or argued for that construction, or any construction at all. APPX3626-3638 (Netlist's claim construction arguments); APPX3685-3693 (Diablo's claim construction arguments).

It is error for the Board to adopt and apply a new construction for the first time in its final written decision. *See SAS Inst., Inc. v. ComplementSoft, LLC*, --- F.3d ---, No. 15-1346, 2016 WL 3213103, at *6-7 (Fed. Cir. June 10, 2016) (vacating Board decision because it *sua sponte* adopted and applied a new construction in its final decision); 5 U.S.C. § 554(b)(3). As this Court has explained, an agency may not “‘change theories in midstream’” without giving the parties meaningful notice of the change and an opportunity to respond. *SAS*, 2016 WL 3213103, at *7 (quoting *Belden Inc. v. Berk-Tek LLC*, 805 F.3d 1064, 1080 (Fed. Cir. 2015)); *see also In re Magnum Oil Tools Int'l, Ltd.*, --- F.3d ---, No. 15-1300, 2016 WL 3974202, at *10 (Fed. Cir. July 25, 2016) (Board must base its decision on arguments “to which the opposing party was given a chance to respond.”). The Board erred by doing exactly that here, and its decision on this ground should be vacated for that reason alone. *Dell Inc. v. Accelleron, LLC*, 818 F.3d 1293, 1301 (Fed. Cir. 2016) (vacating Board decision because Board denied patent owner's procedural rights by relying on argument raised for the first time at the Board hearing).

In any event, the Board’s decision also should be vacated because its belated construction is wrong. The Board did not actually evaluate the claim language (or its incorrect version of the claim language) or the specification. It instead justified its interpretation based “on the same reasoning” it used to justify its construction of “selectively electrically coupling.” APPX64. The Board simply changed “making a selection . . . to transfer power” (from the “selectively electrically coupling” construction) to “making a selection . . . and not transferring power” (in the “selectively isolating/isolate” construction). APPX64.

The Board’s interpretation is wrong for the same reason it was wrong with regard to the identical phrase in the ’536 patent. Although the Board used different words when construing “selectively isolating” in the ’150 and ’536 patents (a difference the Board never explains), it made clear that it interpreted the phrases identically in both patents. APPX93 (explaining its interpretation of the ’150 patent), APPX123 (giving the same explanation for the ’536 patent). The Board’s conclusion of unpatentability for anticipation of claims 22 and 26 of the ’150 patent thus should be vacated.

C. The Board’s Conclusions Of Unpatentability For Obviousness Of All Claims Of The ’150 Patent Fall With Its Erroneous Claim Interpretation Because The Board Evaluated All Claims Identically

The Board’s conclusions of unpatentability for obviousness of claims 22, 24, and 26 of the ’150 patent fare no better. In addition to its flawed anticipation

finding, the Board concluded that these claims were unpatentable because they would have been obvious in light of various combinations of art with Amidi. But unlike its analysis of anticipation, the Board never evaluated the “selectively isolating” limitation of these claims at all. It simply lumped in these claims with the claims requiring “selectively electrically coupling.”

For example, the Board titled its analysis of the Amidi-Klein combination “Amidi and Klein Teach or Suggest All the Recited Limitations of Independent Claims 15, 22, and 31.” APPX23. The Board concluded, among other things, that Amidi disclosed “selectively electrically coupling” because Amidi included using chip-select signals in selecting an active device. APPX29. But the Board never mentioned the separate requirement of “selectively isolating” required by claim 22. APPX28-33. It simply concluded that claim 22 would have been obvious together with the “selectively electrically coupling” claims. APPX33 (“Accordingly, we hold that Petitioner has shown by a preponderance of the evidence that claims 15, 22, and 31 of the ’150 patent are unpatentable . . .”).

The Board did the same when analyzing the combination of Amidi and Wiggers. It stated that its analysis applied to “Claims 15-17, 22, 24, 26, and 31-33.” APPX34. It assessed whether the combination would have included “selectively electrically coupling” under its adopted construction. APPX47-48. But it did not perform any analysis or even mention the “selectively isolating”

limitation of claim 22 and its dependent claims. APPX42-48. It nevertheless held all claims unpatentable. APPX48-49.

And the Board's analysis of the combination of Amidi and Ludwig involved more of the same. APPX76 (lumping together all independent claims). The Board concluded that the combination would have included the "selectively electrically coupling" limitation under its construction, but it failed to analyze, let alone mention, the "selectively isolating" limitation of claims 22, 24, and 26. APPX80-88.

The Board's conclusion of obviousness with regard to claims 22, 24, and 26 thus falls with its flawed obviousness analysis of the '150 patent's other claims. The only analysis the Board gave was tainted by its erroneous interpretation of the claims.

III. THE BOARD UNREASONABLY CONSTRUED "A CIRCUIT CONFIGURED TO BE MOUNTED ON A MEMORY MODULE" TO REACH PART OF A CIRCUIT CONFIGURED TO BE MOUNTED ON PART OF A MEMORY MODULE

The challenged circuit claims of both patents require "a circuit configured to be mounted on a memory module." The Board construed this term to reach "circuitry configured to be mounted on at least a portion of a memory module." APPX13, APPX64, APPX106. That construction is inconsistent with the claim

language and shared description for both patents, and it is unmoored from the purposes of the inventions.⁹

First, the Board’s construction is divorced from the claim language. A “circuit configured to be mounted on a memory module” requires simply what it states—an entire circuit mounted on a single memory module. None of those words is ambiguous. Nor does any suggest that it is enough if only a *portion* of the circuit is mounted on the memory module, with other portions of the circuit mounted elsewhere.

The Board justified its construction by relying on the word “circuit.” A “circuit,” the Board reasoned, is a broad term that covers a variety of electrical components or devices connected in a variety of ways. APPX14, APPX65, APPX107. According to the Board, nothing about the word “circuit” suggests the claims are limited to “only a configuration of electrical components or devices that are mounted on a single memory module.” APPX14, APPX65, APPX107.

The Board’s reasoning does not support its conclusion. The phrase the Board was required to construe is not “circuit.” It is “circuit configured to be mounted on a memory module.” And it is not the word “circuit” in that phrase that Netlist argued limits the claim to circuits mounted on a memory module. It is the

⁹ The challenged circuit claims are claims 15-17, 22, 24, 26, and 31-33 of the ’150 patent and claims 1 and 16-17 of the ’536 patent.

words “mounted on a memory module.” The Board simply stopped its analysis at the word “circuit” and failed to consider the entire claim phrase. The broadest reasonable interpretation standard does not give the Board license to adopt a broad construction by ignoring wholesale specific words in a claim. *Dell*, 818 F.3d at 1299-300.

Second, the Board’s construction defeats an important purpose of the inventions. A key feature of the inventions is increasing the memory capacity of existing computer systems, often at reduced cost compared to alternatives. APPX170-171 (col.14:45-col.15:15). The inventions achieve this by increasing the number of memory devices that can be included on a single memory module. APPX171 (col.15:1-34). The patents teach including a circuit on a memory module that mitigates the properties of existing computer systems that would otherwise limit the number of devices on a module, including the memory controller’s ability to address only a limited number of devices and the performance consequences of increased loading caused by an increased number of devices. APPX164 (col.2:43-59), APPX167 (col.8:48-56); *see also supra*, pp. 41-42. With the patents’ teachings, an existing system’s memory capacity could be increased simply by swapping a conventional memory module for one designed according to the claimed inventions; no other changes to the system would be necessary. APPX2490 (¶63).

By contrast, if a circuit with the claimed features were to have portions mounted on a memory module and other portions mounted elsewhere in a computer system, the intended benefit of easily swapping a higher capacity memory module according to the claimed inventions in place of a conventional memory module would be lost. Instead, the increased memory capacity offered by the inventions could be achieved only by re-wiring a system to add the off-module portions and making other significant hardware alterations. Indeed, if changes to the computer system are possible, the problem the inventions solve no longer exists—overcoming the limitations of a computer system that supports a *fixed, non-changeable* number of memory devices per memory module. Thus, the Board’s interpretation “is unmoored from, rather than aligned with, the description of the invention[s]” and is improper. *See World Class Tech. Corp. v. Ormco Corp.*, 769 F.3d 1120, 1124 (Fed. Cir. 2014).

The Board’s erroneous construction of “circuit configured to be mounted” led it to err at least with regard to unpatentability based on the combinations of Amidi and Wiggers. APPX42-44. The Board concluded that the combination rendered unpatentable claims of the ’150 patent based on Wiggers’ disclosure of switches spanning multiple memory modules. APPX43. The Board thought it immaterial that such a circuit was not mounted on a single memory module: Wiggers disclosure of switches “mounted on memory *modules* controlled by a

single memory controller falls within the scope of the term ‘circuit configured to be mounted on a memory module’ as [the Board] ha[d] construed the term.” APPX43-44. The Court should vacate that conclusion and remand with instructions for the Board to reevaluate the references based on the correct construction.

CONCLUSION

The Board’s final decisions should be vacated and remanded.

Dated: August 22, 2016

Respectfully submitted,

/s/ Brian R. Matsui
BRIAN R. MATSUI
SETH W. LLOYD
MORRISON & FOERSTER LLP
2000 Pennsylvania Avenue NW
Washington, DC 20006
Telephone: (202) 887-8784
BMatsui@mofo.com

MEHRAN ARJOMAND
MORRISON & FOERSTER LLP
707 Wilshire Boulevard
Los Angeles, CA 90017
Telephone: (213) 892-5630

Counsel for Appellant Netlist, Inc.

ADDENDUM

TABLE OF CONTENTS

	Page
Final Written Decision, IPR2014-00882, 12/14/2015.....	APPX1
Final Written Decision, IPR2014-01011, 12/14/2015.....	APPX52
Final Written Decision, IPR2014-00883, 12/14/2015.....	APPX98
U.S. Patent No. 7,881,150.....	APPX136
U.S. Patent No. 8,081,536.....	APPX187

Trials@uspto.gov
571.272.7822

Paper No. 33
Filed: December 14, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DIABLO TECHNOLOGIES, INC.,
Petitioner,

v.

NETLIST, INC.,
Patent Owner.

Case IPR2014-00882
Patent 7,881,150 B2

Before LINDA M. GAUDETTE, BRYAN F. MOORE, and
GEORGIANNA W. BRADEN, *Administrative Patent Judges*.

BRADEN, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318 and 37 C.F.R. § 42.73

IPR2014-00882
Patent 7,881,150 B2

I. INTRODUCTION

We have jurisdiction to hear this *inter partes* review under 35 U.S.C. § 6(c). This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of US Patent No. 7,881,150 B2 (Ex. 1001, “the ‘150 patent”) are unpatentable.

A. Procedural History

Diablo Technologies, Inc. (“Petitioner”) filed a Corrected Petition (Paper 5, “Pet.”) to institute an *inter partes* review of claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent. Netlist, Inc. (“Patent Owner”) filed a Preliminary Response (Paper 9, “Prelim. Resp.”). Pursuant to 35 U.S.C. § 314(a), we instituted an *inter partes* review of all challenged claims on the following grounds alleged in the Petition.

References	Basis	Claims Challenged
Amidi ¹ and Klein ²	§ 103	15–17, 22, 24, 26, and 31–33
Amidi and Wiggers ³	§ 103	15–17, 22, 24, 26, and 31–33

Paper 11 (“Dec. to Inst.”), 33.

After institution of trial, Patent Owner filed a Patent Owner Response (Paper 25, “PO Resp.”), to which Petitioner filed a Reply (Paper 27, “Reply”). An oral argument was held on July 28, 2015, consolidated with

¹ US Patent Publication No. 2006/0117152 A1, pub. June 1, 2006 (filed Jan. 5, 2004) (“Amidi,” Ex. 1008).

² US Patent Publication No. 2001/0008006 A1, pub. July 12, 2001 (“Klein,” Ex. 1009).

³ US Patent No. 6,011,710, iss. Jan. 4, 2000 (“Wiggers,” Ex. 1010).

IPR2014-00882
Patent 7,881,150 B2

the oral hearings for IPR2014-00883 and IPR2014-01011. *See* Paper 30. A transcript (“Tr.”) of the oral argument is included in the record. Paper 31.

B. Related Proceedings

Petitioner informs us that the ’150 patent is involved in the following federal district court cases: *Diablo Technologies, Inc. v. Netlist, Inc.*, Case No. 4:13-CV-03901-YGR (N.D. Cal.); and *Netlist, Inc. v. Smart Modular Technologies*, Case No. 4:13-CV- 05889-YGR (N.D. Cal.). Paper 10, 1. In addition, Petitioner filed two other petitions requesting *inter partes* review of the ’150 patent and related U.S. Patent No. 8,081,536 B1. *Id.* at 2. These cases are: IPR 2014-00883 and IPR2014-01011. *Id.* We consolidated the oral hearings for IPR2014-00882, IPR2014-00883, and IPR 2014-01011. *See* Paper 30.

Petitioner further informs us that related US Patent Nos. 7,619,912 and 7,636,274 are the subjects of *inter partes* reexamination (95/000,578 and 95/001,337). Pet. 10–11. Petitioner also informs us that related U.S. Patent No. 7,289,386 is the subject of district court case *Google, Inc. v. Netlist, Inc.*, Case No. C 08-4144 SBA (N.D. Cal.). *Id.* at 14–15.

C. The ’150 Patent

The ’150 patent relates to a memory module of a computer system with improved performance and memory capacity. Ex. 1001, 1:30–34. Memory module 10 includes a plurality of memory devices 30 (arranged in ranks 32) and circuit 40. *Id.* at 4:56–65; Fig. 1. Circuit 40 is electrically coupled to the memory devices 30 and memory controller 20 of a computer system. *Id.* The memory module improves performance and memory capacity by isolating electrical loads of the memory devices from the computer system. *Id.* at 4:65–66.

IPR2014-00882
Patent 7,881,150 B2

Circuit 40 receives input signals from memory controller 20. *Id.*

Figure 1, reproduced below, illustrates input signals (corresponding to a number of memory devices) from memory controller 20, such as chip select signals (“cs#”), that are directed to memory module 10, which can act as a virtual memory module. *Id.* at 16:47–57; Figs. 1, 9A, 9B.

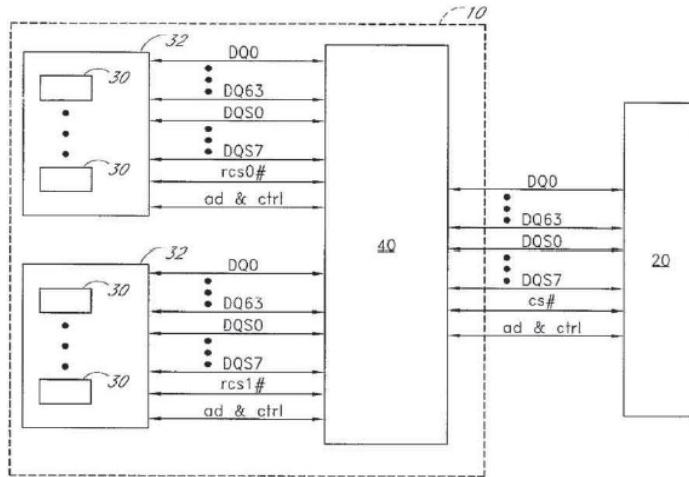


Figure 1 is a schematic of a memory module with circuit 40 and memory devices 30 and connectable to memory controller 20.

As shown in Figure 1 above, based on the received input signals from memory controller 20, circuit 40 generates output signals corresponding to memory devices 30 on the memory module. *Id.* at Fig. 1. The output signals include a different number of chip select signals (e.g., “rcs0#” and “rcs1#”) corresponding to memory devices 30 shown in ranks 32. *Id.* at 16:66–17:4; Fig. 1.

Circuit 40 includes a logic element, such as a programmable logic device (PLD), an application-specific integrated circuit (ASIC), a field programmable gate array (FPGA), or a complex programmable-logic device (CPLD). Ex. 1001, 6:4–18. As shown in Figure 9A, reproduced below, circuit 40 may also include register 230 and phase-lock loop device (PLL) 220. *Id.* at 15:35–41; Fig. 9A.

IPR2014-00882
Patent 7,881,150 B2

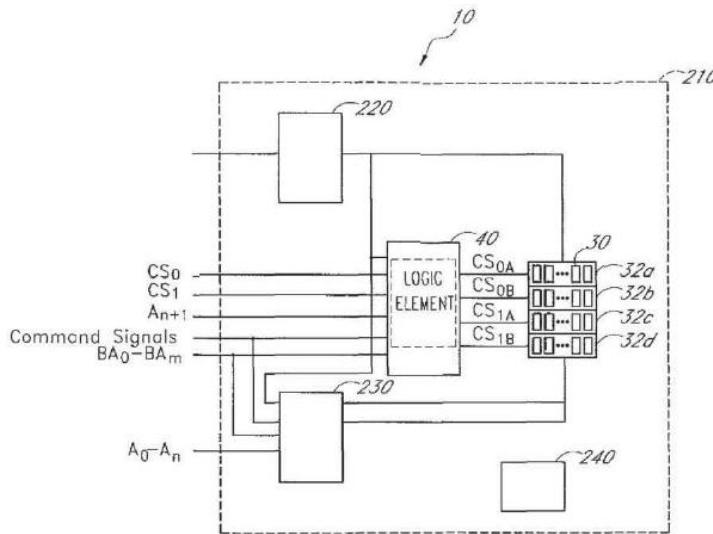


Figure 9A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Figure 9A above illustrates circuit 40 receiving a set of input command signals, address signals (A_{n+1}), including bank address signals (BA_0-BA_m), row address signals (A_0-A_n), column address signals, gated column address strobe signals, and chip-select signals (CS_0, CS_1), from memory controller 20 of the computer system. Ex. 1001, 16:24–29, 17:11–26. In response to the set of input address and command signals, circuit 40 generates a set of output address and command signals. *Id.* at 16:31–33.

With the output address and command signals, circuit 40 isolates the electrical loads of some memory devices 30 from the computer system. *Id.* at 6:48–62. According to the '150 patent, load isolation may result in specific benefits including reduced load related to data signal lines. *Id.* at 14:34–40. In order to isolate the loads, the logic element of circuit 40 translates between a system memory domain of the computer system and a physical memory domain of memory module 10. *Id.* at 6:48–62. As shown in Figure 3, reproduced below, the circuit isolates the load of a memory

IPR2014-00882
Patent 7,881,150 B2

device by isolating one or both of DQ data signal lines 102a, 102b of two memory devices 30a and 30b from common DQ data signal line 112 that is coupled to the computer system. *Id.* at 6:63–7:2, Fig. 3A.

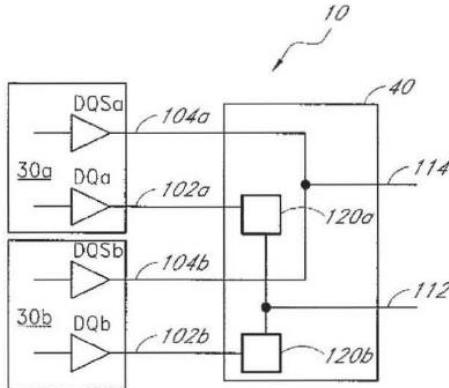


Figure 3A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Circuit 40, shown in Figure 3A above, can electrically couple one or both of DQ data signal lines 102a, 102b of memory devices 30a and 30b to common data signal line 112, at the same time. *Id.* at 7:22–26; Fig. 3A.

Circuit 40 also allows a DQ data signal to be transmitted from memory controller 20 of the computer system to one or both of DQ data signal lines 102a, 102b. *Id.* at 7:2–5. The logic element of circuit 40 uses switches 120a, 120b in order to isolate or couple one or both of DQ data signal lines 102a, 102b of memory devices 30a and 30b from common data signal line 112. *Id.* at 7:2–12.

D. Illustrative Claim

As noted above, *inter partes* review was instituted for claims 15–17, 22, 24, 26, and 31–33 of the '150 patent, of which claims 15, 22, and 31 are independent claims. Claim 15 is illustrative of the challenged claims and is reproduced below:

IPR2014-00882
Patent 7,881,150 B2

15. A circuit configured to be mounted on a memory module so as to be electrically coupled to a first double-data-rate (DDR) memory device having a first data signal line and a first data strobe line, to a second DDR memory device having a second data signal line and a second data strobe line, and to a common data signal line, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;

a register;

a phase-lock loop device configured to be operationally coupled to the first DDR memory device, the second DDR memory device, the logic element, and the register,

wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the memory module, wherein the system memory domain has a first memory density per memory device, and the physical memory domain has a second memory density per memory device less than the first memory density per memory device.

Ex. 1001, 42:41–43:2.

II. DISCUSSION

A. *Claim Construction*

In an *inter partes* review, claim terms in an unexpired patent are interpreted according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b); *see also In re Cuozzo Speed Techs., LLC*, 793 F.3d 1268, 1278–79 (Fed. Cir. 2015) (“Congress implicitly approved the broadest reasonable interpretation standard in enacting the AIA,” and “the standard was properly adopted by

IPR2014-00882
Patent 7,881,150 B2

PTO regulation.”). Under this standard, claim terms generally are given their ordinary and customary meaning, as understood by one of ordinary skill in the art in the context of the patent’s entire written disclosure. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). Yet a “claim term will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.” *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002).

In the Decision to Institute, we construed the terms “Memory Module,” “Circuit Configured to be Mounted on a Memory Module,” and “Selectively Electrically Coupling,” which are recited in all the challenged independent claims. *See* Dec. to Inst. 7–12. During the course of the trial, Patent Owner argued for altered constructions of these claim terms. PO Resp. 5–16. Therefore, we address these contentions and construe each claim term as discussed below.

1. “Memory Module”

In the Decision to Institute, we construed the term “memory module,” as “a plurality of memory devices and a circuit” thereby encompassing “additional circuitry and multiple printed circuit boards.” Dec. to Inst. 7–9.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 4; Tr. 5:25–6:10. Patent Owner, however, contends that “memory module” should be construed as “a packaging arrangement of one or more memory device(s) for use in a computer socket.” PO Resp. 5; Tr. 47:5–7. According to Patent Owner, the construction of “memory module” in the Decision to Institute is unreasonably broad, because it is inconsistent with the ordinary and customary meaning of the term as it would be

IPR2014-00882
Patent 7,881,150 B2

understood by a person of ordinary skill in the art. PO Resp. 6–10. Patent Owner argues the Board erred by construing the term by analyzing each component of the word separately (*id.* at 6) and relying on the '150 patent specification (Tr. 47:17–20), whereas a person of ordinary skill in the art would have understood “memory module” to be a term of art (PO Resp. 6; Ex. 2002 ¶ 54 (Declaration of Dr. Carl Sechen)). Patent Owner explains that under the Board’s construction of “memory module,” the term would encompass a memory controller and associated memory devices. PO Resp. 10.

Patent Owner further contends that the Board’s construction of “memory module” is inconsistent with the '150 patent disclosure. *Id.* at 11. Patent Owner notes that claims 15, 22, and 31 recite “the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals.” *Id.* (emphasis omitted). According to Patent Owner, due to the use of different terms in the '150 patent, “memory module” would not be read as including the '150 patent’s “memory controller” by a person of ordinary skill in the art. *Id.*; Ex. 2002 ¶ 59.

We are charged with interpreting claim terms according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b). Additionally, when construing claim terms, we “should also consult the patent’s prosecution history in proceedings in which the patent has been brought back to the [U.S. Patent and Trademark Office] for a second review.” *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015). Yet, we must be careful not to improperly import limitations into the claims or to read a particular

IPR2014-00882
Patent 7,881,150 B2

embodiment appearing in the written description into the claim, if the claim language is broader than the embodiment. *In re Van Geuns*, 988 F.2d 1181, 1184 (Fed. Cir. 1993).

The specification of the '150 patent does not define explicitly the term "memory module." The specification does, however, teach embodiments that describe a memory module as comprising a plurality of memory devices on a carrier and a circuit. Ex. 1001, 2:63–64; 3:7–9; 4:59–63. In another embodiment, a memory module comprises (i) a printed circuit board on which memory devices are mounted, (ii) a plurality of edge connectors configured to be electrically coupled to a corresponding plurality of contacts of a module slot of the computer system, and (iii) a plurality of electrical conduits which electrically couple the memory devices to the circuit and which electrically couple the circuit to the edge connectors. *Id.* at 5:24–32. The '150 patent also teaches that memory modules in the disclosed embodiments are compatible with at least single in-line memory modules (SIMMS) and dual in-line memory modules (DIMMS). *Id.* at 5:32–39.

Although the embodiments disclosed in the '150 patent are instructive, the claims recite language broader than that found in the embodiments. See *In re Van Geuns*, 988 F.2d at 1184. Therefore, we decline to adopt Patent Owner's claim construction as it would import limitations improperly from the specification into the claims and unnecessarily limit the scope of the claims. We credit, however, the testimony of Patent Owner's Declarant, Dr. Sechen, who explains the state of the art and the customary meaning of "memory module" as it would be understood by one of ordinary skill in the art to encompass at least a "removable circuit board, cartridge, or other carrier that contains one or

IPR2014-00882
Patent 7,881,150 B2

more RAM memory chips.” See Ex. 2002 ¶¶ 41–59. Therefore, we modify the construction of “memory module” from that set forth in the Decision to Institute, wherein we construed the term as “a plurality of memory devices and a circuit” that “encompasses additional circuitry and multiple printed circuit boards.” Dec. to Inst. 9. Rather, we construe the term “memory module” as “one or more memory devices on a carrier,” because such a construction is consistent with the disclosure of the ’150 patent and with the ordinary and customary meaning of “memory module.”

2. “*Selectively Electrically Coupling*”

In the Decision to Institute, we construed the term “selectively electrically coupling,” as “making a selection between at least two components so as to transfer power or signal information from one component to at least one other component.” Dec. to Inst. 9–11.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 4; Tr. 5:25–6:10. Patent Owner, however, contends that “selectively electrically coupling” should be construed as “electrically coupling in response to a selection.” PO Resp. 13–16; Tr. 70:4–9. According to Patent Owner, the Board’s construction is unreasonably broad, whereas its proffered construction is more consistent with the disclosure of the ’150 patent. PO Resp. 15 (citing Ex. 1001, 5:29–30, 7:22; Ex. 2002 ¶¶ 68–69). Patent Owner specifically argues that “electrically coupling” in the ’150 Patent is provided by a structural pathway for electricity, and this is also consistent with the meaning of “electrically coupling” as a term of art. *Id.* (citing Ex. 2002 ¶ 68). Patent Owner further argues that a person of ordinary skill in the art would understand the act of electrically coupling to

IPR2014-00882
Patent 7,881,150 B2

take place between strictly two components, between which a structural pathway for electricity would be formed. *Id.*

We are unpersuaded by Patent Owner's position. The specification of the '150 patent does not define explicitly the term "selectively electrically coupling." Therefore, we refer to its ordinary and customary meaning, as would be understood by one of ordinary skill in the art in the context of the entire disclosure. *In re Translogic Tech., Inc.*, 504 F.3d at 1257 (Fed. Cir. 2007). A technical dictionary, the IEEE Dictionary,⁴ defines "electrical coupling" as "[e]lectrical charges in conductors of a disturbed circuit formed by electrical induction." Ex. 3001. The IEEE Dictionary explains that "[s]ince the ratio of a conductor's electrostatic charge to the potential difference between conductors (required to maintain that charge) is the general definition of capacitance, electrical coupling is also called capacitive coupling." *Id.* The IEEE Dictionary defines "coupling capacitance (1) (ground systems)" ("capacitive coupling") as "[t]he association of two or more circuits with one another by means of capacitance mutual to the circuits." Ex. 3002. We understand this to mean that the two or more circuits are associated in such a way that power or signal information may be transferred from one circuit to another. The Oxford English Dictionary defines "selectively" as "[i]n a selective manner; by selection." Ex. 3003. The Oxford English Dictionary also defines "select" as "[t]o choose or pick out in preference to another or others." Ex. 3004.

⁴ IEEE 100: The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition, Standards Information Network, IEEE Press (2000).

IPR2014-00882
Patent 7,881,150 B2

Accordingly, we modify slightly the construction from the Decision to Institute of “selectively electrically coupling,” as “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component,” because such a construction is consistent with the ordinary and customary meaning of “selectively electrically coupling.”

3. *“Circuit Configured to be Mounted on a Memory Module”*

In the Decision to Institute, we construed the term “a circuit configured to be mounted on a memory module,” to encompass “circuitry configured to be mounted on at least a portion of a memory module.” Dec. to Inst. 11–12. Such a construction is consistent with the ordinary and customary meaning of “a circuit configured to be mounted on a memory module.” *Id.* at 12.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 4; Tr. 5:25–6:10. Patent Owner, however, contends that “a circuit configured to be mounted on a memory module” should be construed as “an entire circuit configured to be mounted on a single memory module.” PO Resp. 11–13; Tr. 68:1–18.

Patent Owner notes that our claim construction, as set forth in the Decision to Institute, is ambiguous in that it can be read two ways:

One might read the Board’s construction as meaning that the circuit is mounted on and occupies at least a portion of the memory module (Ex. 2002, ¶ 62), which may be consistent with Netlist’s construction. On the other hand, one might read the Board’s construction as encompassing portions of the circuit to be mounted off-module, which would be unreasonably broad to a [person of ordinary skill in the art].

IPR2014-00882
Patent 7,881,150 B2

Id. at 11 (citing Ex. 2002 ¶ 62). According to Patent Owner, “memory module” is a term of art that would have had a well-understood meaning to a person of ordinary skill in the art at the time of the invention and a person of ordinary skill in the art would not have understood “a circuit configured to be mounted on a memory module” to include circuit parts off of that memory module or on a different memory module. *Id.* at 12 (citing Ex. 2002 ¶ 60). Patent Owner contends that such a construction could include situations that defeat the purpose of a memory module to make removing and installing memory upgrades easy and error-free. *Id.* (citing Ex. 2002 ¶ 61).

We decline to adopt Patent Owner’s claim construction as it is inconsistent with the definition of “circuit” as found in the specification of the ’150 patent. The ’150 patent defines “circuit” as “a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.” Ex. 1001, 5:9–13. The ’150 patent does not limit a “circuit” to only a configuration of electrical components or devices that are mounted on a single memory module. Therefore, applying the broadest reasonable interpretation consistent with the specification of the ’150 patent, we construe the claim element “a circuit configured to be mounted on a memory module,” as we did in the Decision to Institute, but we further clarify the construction to encompass “at least a portion of circuitry configured to be mounted on at least a portion of a memory module.”

IPR2014-00882
Patent 7,881,150 B2

4. Other Claim Terms

We determine that no express constructions of any other claims terms are required for our analysis, and we apply the ordinary and customary meaning of each claim term.

B. Principles of Law

To prevail in its challenges to the patentability of the claims, a petitioner must establish facts supporting its challenges by a preponderance of the evidence. 35 U.S.C. § 316(e); 37 C.F.R. § 42.1(d). A claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

We analyze the instituted grounds of unpatentability in accordance with the above-stated principles.

C. Level of Ordinary Skill in the Art

In determining whether an invention would have been obvious at the time it was made, we determine the level of ordinary skill in the pertinent art at the time of the invention. *Graham*, 383 U.S. at 17. “The importance of resolving the level of ordinary skill in the art lies in the necessity of

IPR2014-00882
Patent 7,881,150 B2

maintaining objectivity in the obviousness inquiry.” *Ryko Mfg. Co. v. Nu-Star, Inc.*, 950 F.2d 714, 718 (Fed. Cir. 1991).

Petitioner’s Declarant, Srinivasan Jagannathan, Ph.D. (“Dr. Jagannathan”), testifies that a person of ordinary skill in the art at the time of the ’150 patent:

would understand basic memory and data communication concepts, with a bachelor’s degree in any of electrical engineering, computer engineering, computer science, or related field. Course work for one of ordinary skill in the art would have included a course on computer organization, principles of digital design, or computer architecture. One of ordinary skill in the art would also have around one year of experience related to computer memory systems. For example, such experience may include experience in DRAM memory technology and related industry standards such as JEDEC standards for DRAM memories and memory modules.

Ex. 1007 ¶ 53. Patent Owner’ Declarant, Carl Sechen, Ph.D. (“Dr. Sechen”), testifies that one of ordinary skill in the art at the time of the ’150 patent would have had an undergraduate degree in electrical engineering or computer engineering, at least two years of professional experience in the design of memory systems, familiarity with the latest JEDEC standard specifications for memory devices and modules, and familiarity with the latest DRAM memory devices widely available in the market. Ex. 2002 ¶ 14. Dr. Sechen further testifies that a person of ordinary skill in the art would have design proficiency in memory modules comprising DDR memory technology, such as memory modules with JEDEC standard DDR SDRAM devices. *Id.* ¶ 15.

Based on our review of the ’150 patent and the types of problems and solutions described in the ’150 patent and cited prior art, we conclude a

IPR2014-00882
Patent 7,881,150 B2

person of ordinary skill in the art at the time of the '150 patent would have a Bachelor's degree in electrical engineering, computer engineering, computer science, or related field, and at least one year of work experience, including familiarity with computer memory systems and related industry standards such as JEDEC standards for DRAM memories and memory modules. We further note that the applied prior art reflects the appropriate level of skill at the time of the claimed invention. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001).

D. Expert Testimony

Patent Owner argues that Petitioner's Declarant, Dr. Jagannathan, does not qualify as a person of ordinary skill in the art at the time of the invention because of his alleged lack of experience designing memory modules. PO. Resp. 16–17. According to Patent Owner, Dr. Jagannathan's experience and background is directed to software and is not relevant to the case. *Id.* at 18–23. Patent Owner argues that, unlike Dr. Jagannathan, its Declarant Dr. Sechen has significant practical experience designing memory modules. *Id.* at 18; Ex. 2001 ¶¶ 3, 4, Exhibit A.

As to his hardware and memory design experience, Dr. Jagannathan was awarded a Doctorate degree in Computer Science, and has two decades of experience in the design, development, and analysis of a wide range of hardware, software, network and database systems. Ex. 1007 ¶ 2. He has designed and implemented hardware virtual memory caches, and researched theoretical performance measures of various cache coherency protocols. *Id.* Dr. Jagannathan also stated during his deposition that he has experience in the design of memory systems. Ex. 2003, 124:12–126:1. Patent Owner contends, however, that Dr. Jagannathan fails to qualify as a person of

IPR2014-00882
Patent 7,881,150 B2

ordinary skill in the art because he lacks sufficient professional experience physically designing (i.e., “actually putting down a design and saying this is what it would be”) a memory module. PO Resp. 21 (citing Ex. 2003, 125:14–17). We disagree.

To testify as an expert under FRE 702, a person need not be a person of ordinary skill in the art, but rather must be “qualified in the pertinent art.” *Sundance, Inc. v. DeMonte Fabricating Ltd.*, 550 F.3d 1356, 1363–64 (Fed. Cir. 2008); see *SEB S.A. v. Montgomery Ward & Co.*, 594 F.3d 1360, 1372–73 (Fed. Cir. 2010) (upholding a district court’s ruling to allow an expert to provide testimony at trial because the expert “had sufficient relevant technical expertise” and the expert’s “knowledge, skill, experience, training [and] education . . . [wa]s likely to assist the trier of fact to understand the evidence”); *Mytee Prods., Inc. v. Harris Research, Inc.*, 439 F. App’x 882, 886–87 (Fed. Cir. 2011) (non-precedential) (upholding admission of the testimony of an expert who “had experience relevant to the field of the invention,” despite admission that he was not a person of ordinary skill in the art). We find that, although Dr. Jagannathan is less experienced than Dr. Sechen in the area of memory module design, he is qualified sufficiently to testify as an expert witness about memory systems and memory modules.

E. Alleged Obviousness of Claims 15–17, 22, 24, 26, and 31–33 in view of Amidi and Klein

Petitioner alleges claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent are unpatentable under 35 U.S.C. § 103 over the combination of Amidi and Klein. Pet. 22–44. Patent Owner disputes Petitioner’s position, arguing that a person of ordinary skill in the art would not have had reason to combine the references in the manner proposed by Petitioner (PO Resp.

IPR2014-00882
Patent 7,881,150 B2

35) and further that the combination of the references fails to teach or suggest all of the claim limitations (*id.* at 25–35, 44–47, 56–58).

We have reviewed the Petition, the Patent Owner Response, and Petitioner's Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of the '150 patent are unpatentable as obvious over the combination of Amidi and Klein.

1. Overview of Amidi

Amidi discloses a memory interface system with a processor, a memory controller, and a memory module. Ex.1008 ¶¶ 2, 3. According to Amidi, a prior art memory interface system is shown in Figure 1, reproduced below.

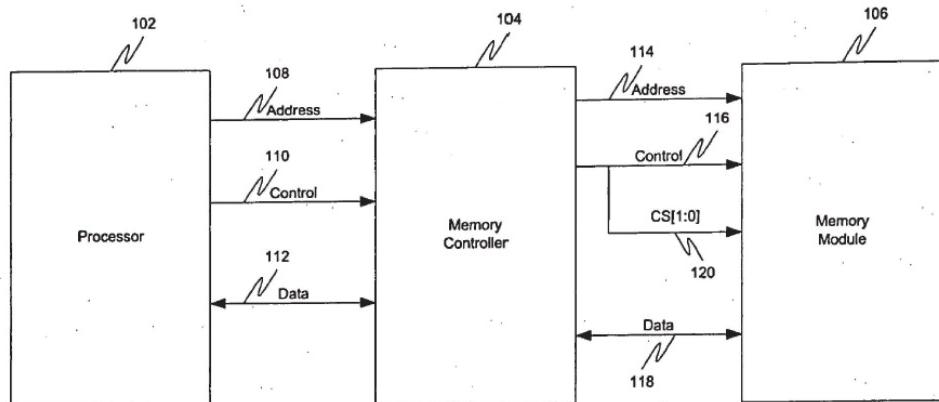


Figure 1 is a schematic of a standard prior art memory interface system.

The prior art system in Figure 1 includes memory module 106 with controller address bus 114, controller control signal bus 116, and controller data bus 118. *Id.* ¶ 2, Fig. 1. As illustrated in Figure 1, memory module 106 communicates with memory controller 104 via busses 114, 116, 118. *Id.* at Fig. 1. Amidi teaches that each stack of DDR memory devices has a data

IPR2014-00882
Patent 7,881,150 B2

signal line and a data strobe line DQS. *Id.* ¶ 32; Fig. 2. Amidi also teaches that at least two DDR memory devices are connected to a common data memory bus. *Id.* ¶ 34; Fig. 3.

Amidi further discloses multiple memory devices mounted on the front and back side of memory module 400 as shown in Figure 4A reproduced below. *Id.* ¶¶ 34, 37.

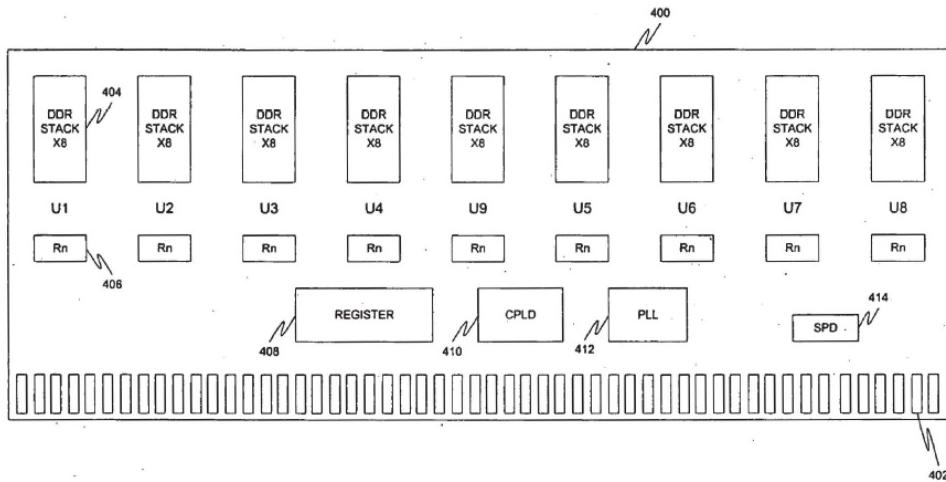


Figure 4A is a schematic of a DDR memory module.

Figure 4A, above, illustrates one embodiment of Amidi where memory module 400 includes memory devices 404, resistor network 406, register 408, complex programmable logic device (CPLD) 410, phase-locked loop (PLL) 412, and SPD 414⁵. *Id.* According to Amidi, memory module 400 receives input signals, including address (Add(n)) signals, row address strobe (RAS) signal, column address strobe (CAS) signal, and bank address (BA[1:0]) signals. Ex. 1008 ¶ 50; Fig. 6A.

⁵ Amidi discloses that SPD 414 is a simple “I2C interface EEPROM [Electrically Erasable Programmable Read-Only Memory] to hold information regarding memory module for BIOS during the power-up sequence.” Ex. 1008 ¶ 40.

IPR2014-00882
Patent 7,881,150 B2

Another embodiment of Amidi's memory interface system is shown in Figure 6A, reproduced below.

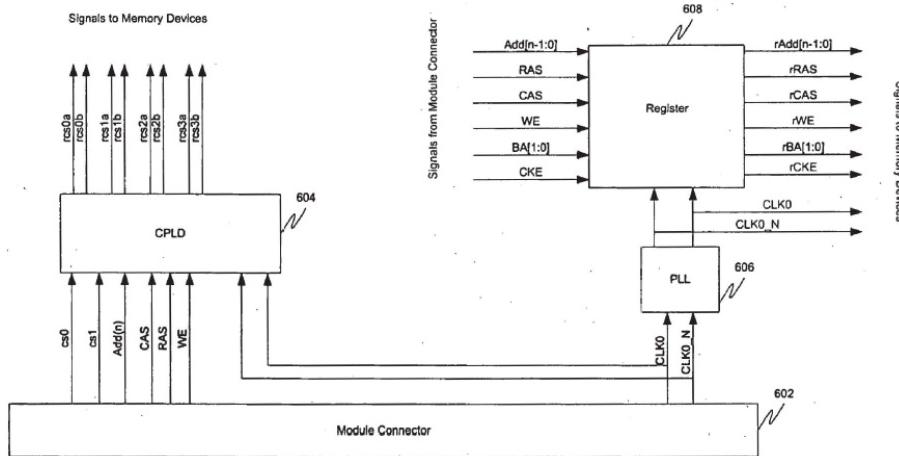


Figure 6A is a schematic of a row address decoding system for a transparent four rank memory module.

As illustrated in Figure 6A above, module connector 602 sends signals to CPLD 604, PLL 606, and register 608. *Id.* CPLD 604 also ensures that all commands for a two rank memory module conveyed by module connector 602 are performed on the four rank memory modules. *Id.* ¶ 52. Amidi explains that the system chip select signals control the ranks of individual memory modules. *Id.* ¶¶ 2, 3.

2. Overview of Klein

Klein discloses a method for bus capacitance reduction. Ex. 1009, Abstract. According to Klein, data bus capacitance is reduced by decoupling unaccessed memory circuits from a data bus during data transfers to or from other memory circuits. *Id.* One embodiment in Klein provides memory controller 22 connects to circuitry 26 for interfacing with one or more memory circuits 28, as shown in Figure 3, reproduced below. *Id.* ¶ 28; Fig. 3.

IPR2014-00882
Patent 7,881,150 B2

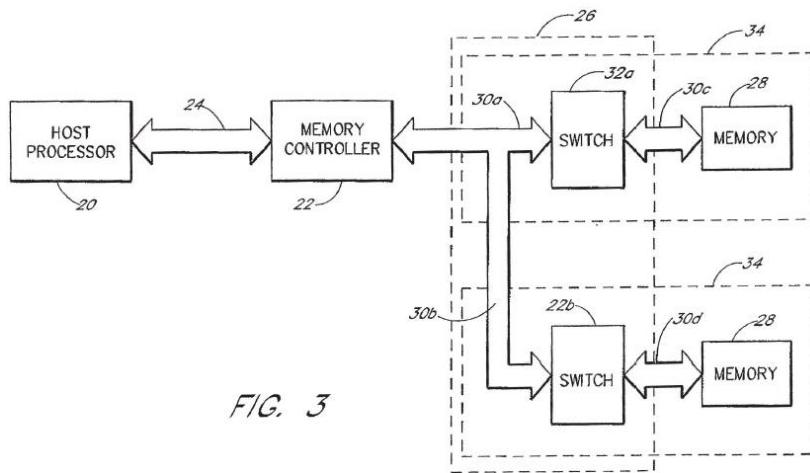


Figure 3 is a schematic of a bus switch that couples or decouples memory elements 28 and memory controller 22.

Figure 3, above, illustrates that the data bus between memory controller 22 and memory elements 28 may comprise several branches 30a, 30b, one for each separate memory elements 28. *Id.* Each branch may include switch 32a, 22b that may be used to selectively isolate portions (30c, 30d) of the data bus running from memory controller 22 to memory circuitry 28. *Id.* Klein states that memory circuit 28 may be a conventional DRAM integrated circuit. *Id.* ¶ 29. According to Klein, the embodiment shown in Figure 3 may reduce the parasitic capacitance that the memory controller needs to charge and discharge during data transfers because a portion of the data bus and the stray capacitance of unaccessed memory circuits are removed. *Id.* ¶ 28.

Another embodiment in Klein is illustrated in Figure 6, reproduced below.

IPR2014-00882
Patent 7,881,150 B2

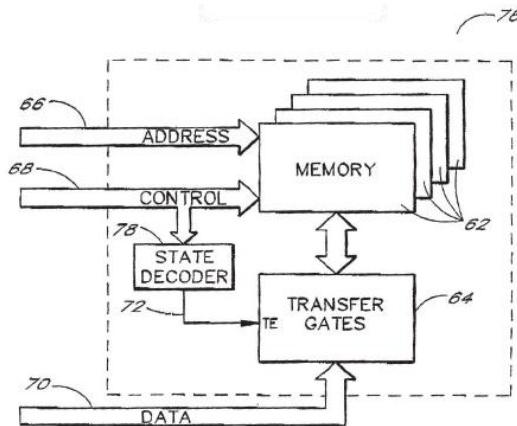


Figure 6 is a schematic of a memory module with memory elements that connect to an integrated circuit with transfer gates and state decoder.

As shown in Figure 6, a circuit is provided on memory module 76 that includes transfer gates 64 and state decoder 78. Ex. 1009 ¶¶ 35, 39. Klein discloses that state decoder 78 includes inverter 80 (*Id.* ¶ 36), and that “the state decoder 78 could comprise a state machine 84 made with a programmable gate array for example” (*Id.* ¶ 37). Also, Klein teaches that the state decoder may be implemented as a state machine. *Id.* ¶ 37, Fig. 8.

Klein further discloses control logic circuitry, data buffer registers, and a bus switch are incorporated into memory modules. *Id.* ¶¶ 29, 39, 40; Figs. 3, 10. According to Klein, the integrated circuit and a transfer gate output are connected to data buffer registers. *Id.* ¶ 40; Fig. 10.

3. Analysis

a. Amidi and Klein Teach or Suggest All the Recited Limitations of Independent Claims 15, 22, and 31

Petitioner contends the combined disclosures of Amidi and Klein, as summarized above, teach or suggest each limitation of independent claims 15, 22, and 31 of the '150 patent. Pet. 22–44. Petitioner first argues that Amidi discloses a circuit mounted on a memory module, where the circuit

IPR2014-00882
Patent 7,881,150 B2

includes a logic element, a register, and a phase-lock loop device. *Id.* at 23 (citing Ex. 1008 ¶¶ 2, 37; Figs. 4A, 6A); Ex. 1007 ¶¶ 58, 64; Reply 12. According to Petitioner, the system described in Amidi includes a memory module having one or more ranks of double-data-rate (DDR) memory devices, which are electrically coupled to the components of the circuit (CPLD). Pet. 24–25 (citing Ex. 1008, Figs. 4A, 4B, 6A); Reply 12–13. Petitioner then explains that Klein also discloses a circuit that is mounted on a memory module and includes a state decoder that may comprise a programmable logic device. Pet. 23–24 (citing Ex. 1009 ¶ 35). Petitioner contends that Klein’s disclosure of circuitry interfacing with memory circuits is a disclosure of a circuit electrically coupled to DDR memory devices. *Id.* at 25.

Petitioner then argues that Amidi discloses the following claim limitations: (i) memory devices having data signal lines and data strobe lines (Ex. 1008 ¶¶ 29, 32; Figs. 2, 3); (ii) stacks of DDR memory devices having a data signal line and a data strobe line DQS (*id.* ¶ 32; Fig. 3); and (iii) at least two DDR memory devices connected to the same (common) memory bus (“common data signal line”) (*id.* ¶¶ 34–35; Fig. 3). Pet. 27. According to Petitioner, a person of ordinary skill in the art would recognize that each DDR memory device has its own data bus and that they are connected to a common data signal line, and therefore, the circuit of Amidi is “electrically coupled” to the common data bus. *Id.* (citing Ex. 1009 ¶¶ 61, 63, 65, 72); Reply 13–14. Petitioner further argues that Klein discloses that bus switch 33 is electrically coupled to output data buses 31c, 31d, 31e, 31f and to a single input data bus 31a (i.e., the “common data signal line”). Pet. 27–28 (citing Ex. 1009 ¶¶ 28, 29); Ex. 1007 ¶¶ 80, 83.

IPR2014-00882
Patent 7,881,150 B2

Petitioner provides arguments that the sending and receiving of input signals by the systems occurs in Amidi and Klein, and explains how each reference teaches a circuit that is responsive to such input signals. Pet. 28–43. Petitioner specifically argues that Amidi teaches a circuit that is “responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line.” Petitioner makes this argument because Amidi teaches that CPLD 404, 604 is responsive to a set of input signals (address signal Add(b) and chip select signals cs0, cs1) to determine (“selectively electrically coupling”) an active rank of the four ranks while inactivating the other three ranks of memory devices from the computer system. *Id.* at 33 (citing Ex. 1008 ¶¶ 43, 44, 62); Ex. 1007 ¶ 72). Petitioner’s Declarant, Dr. Jagannathan, in the Supplemental Declaration (“Ex. 1023”), testifies that when Amidi’s CPLD provides a chip select signal to a rank of memory devices, the signal selects the rank and thereby causes the rank to be coupled to the data bus. Ex. 1023 ¶ 32. According to Petitioner, the act in Amidi of activating one rank and sending a chip select signal to that rank while inactivating other ranks constitutes “selectively coupling” and “selectively isolating.” Pet. 33–36. Petitioner also argues that Klein teaches “selectively isolating,” because Klein discloses that bus switches 32a and 32b (“the circuit”) respond to control lines 68 (“the set of input signals”) to disconnect (“selectively isolating”) one of two output data buses 30c and 30d (“one or more loads of the DDR memory devices”) from input data bus 30a. *Id.* at 36; Reply 16–17.

IPR2014-00882
Patent 7,881,150 B2

According to Petitioner, a person of ordinary skill in the art would have been motivated to combine the teachings of Amidi and Klein, because (1) both references relate to memory devices, such as DIMMS, (2) both references describe coupling or isolating memory device loads, and (3) the combined teachings would result in the benefit of isolating a memory device load from a computer system so as to reduce parasitic capacitance and increase the speed at which memory accesses can be performed. Pet. 43–44 (citing Ex. 1008 ¶¶ 38, 39, 43, 44, 62; Ex. 1009 ¶¶ 9, 10,⁶ 28); Ex. 1007 ¶¶ 89, 91.

Petitioner supports its position with the Declaration of Dr. Jagannathan, who testifies that a person of ordinary skill in the art would have reason to apply the bus switch of Klein to the circuit architecture of Amidi in order to reduce the load seen by the memory controller. Ex. 1007 ¶ 91. Dr. Jagannathan further opines that one of ordinary skill implementing the teachings of Klein would understand that using a circuit that allows for emulating a higher memory density configuration with lower memory density devices provides the predictable benefit of a cheaper implementation as taught by Amidi. *Id.* Dr. Jagannathan notes that Klein specifically teaches how to control the timing of a data bus switch (e.g., the transfer gate switch in Figure 1) and teaches a variety of ways for generating control signals for data bus switches. Ex. 1023 ¶¶ 40–43, 46, 47. According to Dr. Jagannathan, a person of ordinary skill in the art seeking to implement the bus switch of Klein in the memory module of Amidi would have understood

⁶ The Petitioner cites to Ex. 1009, 1:33–2:18, which appears to correspond to paragraphs 9 and 10 of Klein.

IPR2014-00882
Patent 7,881,150 B2

the timing relationships between the chip select signal and when the data is communicated over the data signal lines, because such relationships are dictated by the JEDEC standards (which is cited by Amidi). *Id.* ¶ 45; *see Ex. 1008* ¶ 7.

Patent Owner contests Petitioner's position, arguing that the combination of Amidi and Klein fails to teach or suggest all the recited claim limitations and that a person of ordinary skill in the art would not have had a reason to combine the disclosures of the cited references. PO Resp. 25–32, 35–47.

Patent Owner first contends that the combination of Amidi and Klein fails to teach or suggest a “circuit configured to be mounted on a memory module[,] . . . the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal to the common data signal line,” as recited in the challenged claims. *Id.* at 25–28. Specifically, Patent Owner contends that Petitioner misapplies the prior art because claims 15 and 31 require a “circuit” to perform the “selectively electrically coupling” and the “circuit” is distinct from the claimed “DDR memory device.” *Id.* at 27. Thus, according to Patent Owner, any internal switching function “inside the memory device,” does not meet the claimed “circuit” performing the “selectively electrically coupling.” *Id.*

Patent Owner further contends that Amidi’s disclosure of hard-wiring of all four memory ranks to a data bus constitutes permanent coupling, so that choosing a rank of memory devices does not alter the coupling between data signal lines. *Id.* at 29. Patent Owner argues that direct hard-wiring is static and permanent, and does not respond to a selection. *Id.* at 29–30.

IPR2014-00882
Patent 7,881,150 B2

Patent Owner adds to this argument by contending that “electrically couple” is provided by structure so there is a pathway for electricity, which is different from Amidi’s signal transmission that directs signals down a specific signal line or data bus and constitutes a transmitted flow of electricity. PO Resp. 32 (citing Ex. 1001, 5:29–30; Ex. 2002 ¶ 82).

According to Patent Owner, “[a] pathway for electricity (as for “selectively electrically coupling” in the ’150 Patent) is not a flow of electricity (as for Amidi’s ‘directing signals’).” *Id.* Patent Owner, thus, concludes that Amidi fails to meet the challenged claim limitations because hard-wiring between the memory ranks and the data bus is a permanent coupling and cannot be “selectively electrically coupling.” *Id.* at 30–31.

We do not agree with Patent Owner. Rather, based on the definition of circuit discussed above (*see* Section II.A.3) and as supported by the ’150 patent (*see* Ex. 1001, 5:9–13) and , we find that the combined disclosures of Amidi and Klein teach a “circuit . . . mounted on a memory module” that performs the “selectively electrically coupling” as recited by the challenged independent claims. Specifically, we find that Klein teaches a memory circuit that may be a conventional Dynamic Random Access (“DRAM”) integrated circuit (“IC”) and the DRAM IC may be part of a memory module that also incorporates a separate IC forming the bus switch. Ex. 1009 ¶ 29, Fig. 3. Thus, Klein teaches that the components of a circuit can be installed in a single integrated circuit and mounted on one memory module. *Id.* Therefore, Klein’s disclosure of a memory circuit meets the limitation of a “circuit . . . mounted on a memory module” that provides “selectively electrically coupling.”

IPR2014-00882
Patent 7,881,150 B2

Additionally, we are unpersuaded by Patent Owner’s argument that hard-wired data signal lines, such as that taught by Amidi, cannot be electrically coupled in a selective fashion. As discussed above, we construe “selectively electrically coupling” as “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component.” *See supra* Section II.A.2. Amidi’s disclosure of directing signals down a specific signal line or data bus in order to determine an active rank within the memory devices falls within the scope of the term “selectively electrically coupling” as we have construed the term. Furthermore, Patent Owner’s Declarant, Dr. Sechen, testified that Klein discloses the use of MOSFET switches (i.e., data bus switches) for decoupling select memory circuits from the data bus. *See* Ex. 1022, 49:16–50:18, 88:5–22.

Patent Owner also contends that the combination of Amidi and Klein is improper and based on impermissible hindsight because Amidi’s CPLD output CS signal and Klein’s TE (“transfer enable”) signal are not equivalent due to difference in respective timing operation and purposes. PO Resp. 35–46; *see also* Ex. 1009 ¶ 26 (explaining that transistors are turned on by asserting the gates 15 via an input “transfer enable” signal line labeled TE in Figure 2 and that bus switch circuits such as that illustrated in Figure 2 are known to those of skill in the art.). Patent Owner argues that the specific implementation details, e.g., timing considerations, are critical, and a person of ordinary skill would have found an Amidi-Klein combination inoperable due to timing problems. *Id.* at 38 (citing Ex. 2002 ¶ 127). Patent Owner relies on the Declaration of Dr. Sechen to support its position regarding the

IPR2014-00882
Patent 7,881,150 B2

inoperability of an Amidi-Klein combination. *See* Ex. 2002 ¶¶ 127–130.

Dr. Sechen specifically states that:

due to standardized DDR memory device operation, a DDR read (or write) command’s chip-select signal (e.g., Amidi’s CPLD output chip-select signal) does not coincide with its read (or write) data (e.g., to be gated by Klein’s TE signal). Thus, by failing to properly transfer the target data, the Petitioner’s proposal—DDR chip-select signals from Amidi’s CPLD = Klein’s input TE signal—would malfunction and be inoperable.

Id. ¶ 127.

Patent Owner then argues that an Amidi-Klein combination is also inoperable because a DDR chip-select signal (as used in Amidi) is not designed to be a timing signal, whereas a timing signal is required for Klein’s transfer gates to operate properly. PO Resp. 38–39. Patent Owner supports its position with the Declaration of Dr. Sechen, who testifies that “[f]undamentally, a DDR chip-select signal lacks enough timing information to indicate when to properly open and close Klein’s transfer gates. Thus, a DDR chip-select signal, as output by Amidi’s CPLD, would not be a usable signal at all for controlling Klein’s transfer gates.” Ex. 2002 ¶ 128.

Again, we do not agree with Patent Owner. Rather, we agree with Petitioner’s position and we find that the timing requirements for memory devices are dictated by the JEDEC standards, which were known to a person of ordinary skill in the art at the time of the ‘150 patent. *See e.g.*, Ex. 1008 ¶ 7. Additionally, Klein teaches both (i) the coordination of the signal between the chip-select signal and when the data is being received by the memory device, and (ii) to adjust the timing of the chip-select signals to control access to the device. *See* Ex. 1009 ¶¶ 22–23, Fig. 1. We also credit the testimony of Dr. Jagannathan, who testifies regarding the JEDEC21C-

IPR2014-00882
Patent 7,881,150 B2

4.5.7 standard as it relates to the 168 Pin Registered SDRAM DIMM Family of memory devices. *See e.g.*, Ex. 1007 ¶¶ 43–51. Dr. Jagannathan specifically testifies that

There is a Phase-Locked Loop (PLL) clock input provided to the register and to the memory devices. This is depicted in JEDEC21C-4.5.7, Figure as “PCK” input to the register, and as “CK0 _ PLL” in the case of the memory devices. A phase-locked loop device receives an input clock signal and generates another clock signal whose phase matches (within tolerance) the input clock. The details of the PLL are specified in JEDEC21C-4.5.7 at p. 4.5.7-8. Specifically, the PLL clock output is depicted as driving a number of SDRAM devices and registers.

Id. ¶ 45.

The concept of a PLL would generally be well understood by one of ordinary skill in the art. For instance, Jacob teaches that “[t]he function of a PLL or DLL, in general, is to synchronize two periodic signals so that a certain fixed amount of phase-shift or apparent delay exists between them. The two are similar, and the terms are often used interchangeably.

Id. ¶ 46 (citing Ex. 1018, 11).

Thus, we find that a person of ordinary skill in the art would have known how to address the timing and use of chip-select signals so that the teachings of Klein would have been applicable to Amidi.

Moreover, the test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference. *In re Keller*, 642 F.2d 413, 425 (CCPA 1981). Rather, the test is what the combined teachings of those references would have suggested to those of ordinary skill in the art. *Id.* We credit the testimony of Dr. Jagannathan, who states that “when Klein refers to a ‘signal’ used in any

IPR2014-00882
Patent 7,881,150 B2

of the circuits taught therein, one of ordinary skill would understand it is ‘a varying electrical impulse that conveys information from one point to another.’” Ex. 1007 ¶ 77; *see* Ex. 1023 ¶¶ 33–35. Thus, we are not persuaded that Amidi’s CS signal and Klein’s TE signal are not equivalent in their signaling function, or that the teachings of Amidi and Klein would not have been combinable to one of ordinary skill in the art.

Based on the evidence of record, we agree with Petitioner’s position that challenged claims 15–17, 22, 24, 26, and 31–33 would have been obvious over Amidi and Klein. First, we are persuaded by Petitioner’s reasoning and the evidentiary record that Amidi teaches a circuit mounted on a memory module that is electrically coupled to a first DDR memory device and a second DDR memory device. We are further persuaded that Amidi teaches a circuit with a logic element, a register, and a PLL. We also are persuaded that the teachings of Amidi could have been implemented using the common data signal line and switch system disclosed in Klein so that by (i) selectively electrically coupling a first data signal line to the common data signal line and (ii) selectively electrically coupling a second data signal line to the common data signal line, the circuit is responsive to a set of input signals. Additionally, we credit the testimony of Dr. Jagannathan that a person of ordinary skill in the art would have had a reason to combine the teachings of Amidi with Klein, which both relate to memory devices, such as DIMMS, and describe coupling or isolating memory device loads. *See* Ex. 1007 ¶¶ 89, 91; Ex. 1023 ¶ 45.

Second, the arguments presented by Patent Owner generally attack the references individually, rather than in combination. PO Resp. 25–32. Nonobviousness cannot be established by attacking the references

IPR2014-00882
Patent 7,881,150 B2

individually when a challenge is predicated upon a combination of prior art disclosures. *See In re Merck & Co., Inc.*, 800 F.2d 1091, 1097 (Fed. Cir. 1986); *cf. Keller*, 642 F.2d at 426 (“[O]ne cannot show non-obviousness by attacking references individually where . . . the rejections are based on combinations of references.”). In attacking the references individually, Patent Owner again fails to address Petitioner’s actual challenges and establish an insufficiency in the combined teachings of the references and show Petitioner has not meet its burden in arguing obviousness of the challenged claims.

Lastly, we note that the testimony of Patent Owner’s Declarant, Dr. Sechen, is based on the claim constructions proffered by Patent Owner and not on the constructions set forth in the Decision to Institute. *Compare Ex. 2002 ¶¶ 38–72, with Dec. to Inst. 7–12.* We have considered as relevant, however, the portions of his analysis regarding the prior art and alleged non-obviousness of the claims and accorded the appropriate weight to that particular testimony.

Accordingly, we hold that Petitioner has shown by a preponderance of the evidence that claims 15, 22, and 31 of the ’150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Amidi and Klein.

b. Amidi and Klein Teach or Suggest All the Recited Limitations of Dependent Claims 16, 17, 24, 26, 32, and 33

Claims 16, 24, and 32 recite “wherein the two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.” Ex. 1001, 43:3–5, 44:1–3, 44:58–60. Dependent claim 17 recites that the circuit includes “one or more switches selectively electrically

IPR2014-00882
Patent 7,881,150 B2

coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the one or more switches operatively coupled to the logic element to receive control signals from the logic element.” Dependent claim 33 recites a similar limitation. *Id.* at 43:6–12. Dependent claim 26 further recites that the claimed circuit is “configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.” *Id.* at 44:8–11.

Petitioner contends the combined disclosures of Amidi and Klein, as summarized above, teach or suggest each limitation of dependent claims 16, 17, 24, 26, 32, and 33 of the ’150 patent. Pet. 22–44. Patent Owner does not provide separate contentions regarding additional limitations recited in the dependent claims. *See generally* PO Resp.

After consideration of the language recited in claims 16, 17, 24, 26, 32, and 33 of the ’150 patent, the Petition, the Patent Owner Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers, we find that one of ordinary skill in the art would have considered these dependent claims obvious over the combination of Amidi and Klein. Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that claims 16, 17, 24, 26, 32, and 33 of the ’150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Amidi and Klein.

F. Asserted Obviousness of Claims 15–17, 22, 24, 26, and 31–33 in view of Amidi and Wiggers

Petitioner contends claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent are unpatentable under 35 U.S.C. § 103 in view of Amidi and

IPR2014-00882
Patent 7,881,150 B2

Wiggers. Pet. 44–57. Patent Owner disputes Petitioner’s position, arguing that a person of ordinary skill in the art would not have had reason to combine the references in the manner proposed by Petitioner (PO Resp. 47–58) and further that the combination of the references fails to teach or suggest all of the claim limitations (*id.* at 33–35).

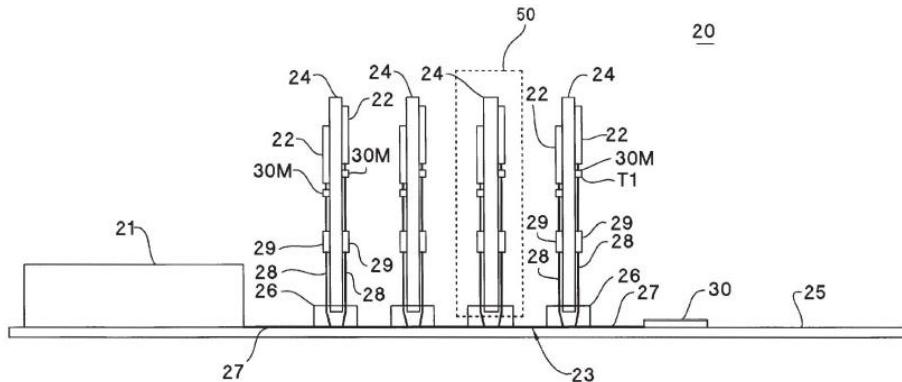
We have reviewed the Petition, the Patent Owner Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent are unpatentable as obvious over the combination of Amidi and Wiggers.

1. Overview of Amidi

See Section II.D.1. discussed above.

2. Overview of Wiggers

Wiggers discloses a system for a reduced capacitance memory system and increased propagation speed for data traveling both from a memory chip to a memory controller and in the reverse direction. Ex. 1010, 1:7–11; 3:17–19. One embodiment of the system disclosed in Wiggers is shown in Figure 3, reproduced below.



IPR2014-00882
Patent 7,881,150 B2

Figure 3 is a schematic of a memory system with memory controller 21, data bus 23, and memory devices 22.

As shown in Figure 3, a central processing unit (CPU) acts as dedicated memory controller 21 that is connected to data bus 23. *Id.* at 4:38–47. Memory controller 21 selectively accesses numerous memory devices 22 which are arranged either serially, in parallel, or in some combination of the two along the data bus 23. *Id.* at 4:47–50. According to Wiggers, the memory devices may include read only memory (ROM) or random access memory (RAM), or dynamic random access memory (DRAM). *Id.* at 4:50–53.

Wiggers further discloses memory devices 22 and switches 29 are preferably affixed to removable memory modules 24 that allow the memory system configuration to be easily changed by simply adding modules or by replacing some or all of the modules. Ex. 1010, 4:61–65. Further, Wiggers discloses switches 29 including field effect transistor (FET) type switches. *Id.* at 4:47–57. Switches 29 of Wiggers are electrically coupled to memory devices 22. *Id.* at 4:53–57.

An embodiment of memory module 24 is shown in Figure 5, reproduced below.

IPR2014-00882
Patent 7,881,150 B2

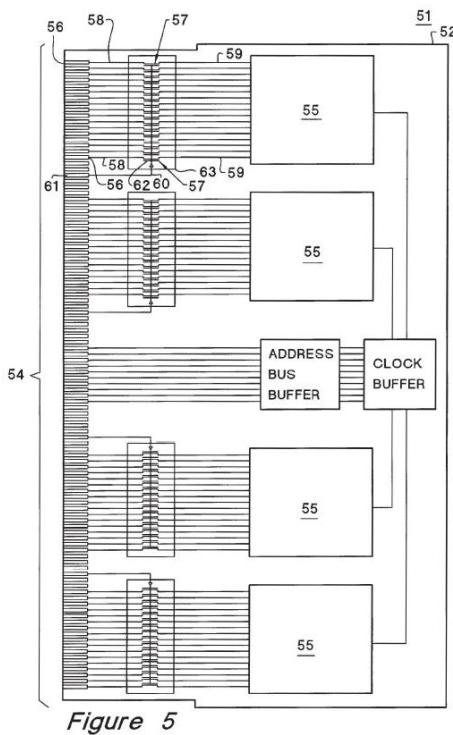


Figure 5

Figure 5 is a detailed plan view of memory module 51, which corresponds to memory module 24 from Figure 3.

Wiggers discloses that Figure 5, above, illustrates “a memory module for reducing the capacitive load in the data bus of a memory system according to the invention and of a type shown in box 50 of Fig. 3.” *Id.* at 6:33–36. Memory module 51 includes substrate 52, at least one memory chip 55 and switches 57 affixed to substrate 52. *Id.* at 6:35–40. Primary data lines 58 connect each data pin to an associated switch and secondary data lines 59 connect each switch to a memory device. *Id.* at 6:41–43. The switches include position controllers 62 for switching the switches between an open and closed position. Ex. 1010, 6:43–45. The position controllers are electrically connected to control line 60 that also electrically connected to control pin 61. *Id.* at 6:46–49.

Another embodiment of the system disclosed in Wiggers is shown in Figure 4, reproduced below.

IPR2014-00882
Patent 7,881,150 B2

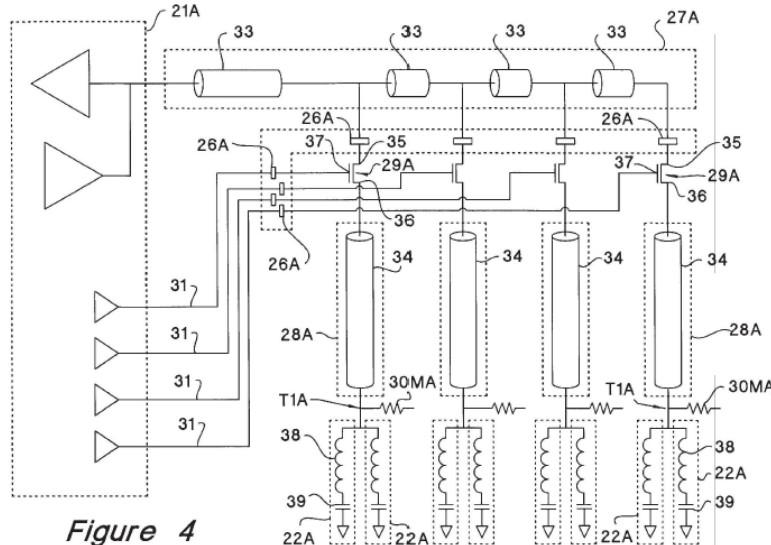


Figure 4

Figure 4 is a schematic of a memory system with the data bus 23 depicted as transmission lines 33.

As shown in Figure 4, data bus 23 (from Figure 3) is illustrated as a series of transmission lines. *Id.* at 5:20–21. Specifically, Wiggers teaches that board portion 27A includes transmission lines 33. *Id.* at 5:21–23.

Wiggers also teaches that the memory controller uses switches to selectively (i) couple a memory device to a data bus when accessing a memory location in the memory device and (ii) decouple the memory device from the data bus at other times. *Id.* at 3:25–28. According to Wiggers, the selective coupling of the memory devices minimizes capacitive loading of the data bus. *Id.* at 3:28–31. Wiggers specifically discloses “there is at least one switch 29A for each data line connected to a memory module.” Ex. 1010, 5:40–41.

Wiggers further discloses that “[w]hen multiple data lines are involved, the switches 29A can be grouped together into set (not shown in Fig. 4) with each switch in the set controlled by a common control signal on a single control line.” *Id.* at 5:47–51.

IPR2014-00882
Patent 7,881,150 B2

3. Analysis

a. Amidi and Wiggers Teach or Suggest All the Recited Limitations of Independent Claims 15, 22, and 31

Petitioner contends the combined disclosures of Amidi and Wiggers, as summarized above, teach or suggest each limitation of independent claims 15, 22, and 31 of the '150 patent. Pet. 44–57. As discussed below, Petitioner presents arguments identifying each claim element in the disclosures of Amidi and Wiggers. Petitioner first argues that Amidi discloses a circuit mounted on a memory module, where the circuit includes a logic element, a register, and a phase-lock loop device. *Id.* at 44 (referencing Ex. 1008 ¶¶ 2, 37; Figs. 4A, 6A); Ex. 1007 ¶ 93. According to Petitioner, the system described in Amidi includes a memory module having one or more ranks of double-data-rate (DDR) memory devices, which are electrically coupled to the components of the circuit (CPLD). *Id.* at 44–45 (citing Ex. 1008, Figs. 4A.) Petitioner then explains that Wiggers discloses a circuit electrically coupled to memory devices. *Id.* (citing Ex. 1010, 4:53–57).

Petitioner argues that Amidi discloses the following claim limitations: (i) memory devices having data signal lines and data strobe lines (Ex. 1008 ¶¶ 29, 32; Figs. 2, 3); (ii) stacks of DDR memory devices having a data signal line and a data strobe line DQS (*id.* ¶ 32; Fig. 3); and (iii) at least two DDR memory devices connected to the same (common) memory bus (“common data signal line”) (*id.* 34–35; Fig. 3). Pet. 27. According to Petitioner, a person of ordinary skill in the art would have recognized that each DDR memory device has its own data bus and that they are connected to a common data signal line, and that the circuit of Amidi is “electrically

IPR2014-00882
Patent 7,881,150 B2

coupled” to the common data bus. *Id.* at 27 (citing Ex. 1009 ¶¶ 61, 63, 65, 72. Petitioner further argues that Wiggers discloses that switches 29A are electrically coupled to a memory device by module portion 28A. *Id.* at 45 (citing Ex. 1010, 6:23–27; Fig. 4); Ex. 1007 ¶ 94. According to Petitioner, given the combined teachings of Amidi and Wiggers, a person of ordinary skill in the art would have found it obvious to organize the memory devices 22 of Wiggers into multiple ranks. Pet. 46 (citing Ex. 1007 ¶ 96). Petitioner then concludes that it would have been obvious to a person of ordinary skill in the art to form a switch that is electrically coupled to a plurality of memory devices arranged in ranks on the memory module, as required by independent claim 22. *Id.*

Petitioner provides arguments for the sending and receiving of input signals by the systems in Amidi and Wiggers, and how each reference teaches a circuit that is responsive to such input signals. *Id.* at 47–52. Petitioner specifically argues that Amidi teaches a circuit that is “responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line,” because Amidi teaches that CPLD 404, 604 is responsive to a set of input signals (address signal Add(b) and chip select signals cs0, csl) to determine (“selectively electrically coupling”) an active rank of the four ranks and inactivating the other three ranks of memory devices from the computer system. Pet. 49 (citing Section V.A.1 of the Petition). According to Petitioner, the act in Amidi of activating one rank while inactivating other ranks constitutes “selectively coupling” and “selectively isolating.” *Id.* Petitioner also argues that Wiggers teaches selectively electrically coupling

IPR2014-00882
Patent 7,881,150 B2

and decoupling individual memory devices (“first data signal line” and “second data signal line”) from the data bus (“common data signal line”).

Id. Petitioner explains that the memory controller can also selectively electrically couple and decouple each of the memory devices from the data bus, either individually or in small groups using a number of switches 29, preferably including field effect transistor (FET) type switches. *Id.* (citing Ex. 1010, 4:53–57). Petitioner further notes that Wiggers discloses that switch 29A (“the circuit”) responds to a control line 31 (“the set of input signals”) from the memory controller 21A to electrically couple or decouple a memory module from the data bus. *Id.* at 49–50 (citing Ex. 1010, 5:40–47).

According to Petitioner, a person of ordinary skill in the art would have been motivated to combine the teachings of Amidi and Wiggers because (1) both references relate to memory devices, (2) both references describe coupling or isolating memory device loads, and (3) the combined teachings would result in the benefit of isolating a memory device load from a computer system so as to reduce parasitic capacitance and increase the speed of data propagation. Pet. 56–57 (citing Ex. 1008 ¶¶ 38, 39, 43, 44, 62; Ex. 1010, 4:27–37; 5:62–66); Ex. 1007 ¶¶ 96.

Petitioner supports its position with the Declaration of Dr. Jagannathan, who testifies that a person of ordinary skill in the art would have had a reason to apply the switches of Wiggers to the memory module of Amidi in order to increase the speed of data propagation. Ex. 1007 ¶ 96 (citing Ex. 1010, 3:39–40; 4:27–37). Dr. Jagannathan further opines that implementing the switches of Wiggers into the architecture of Amidi would

IPR2014-00882
Patent 7,881,150 B2

have been routine for one of ordinary skill in the art and the result of such application would have been as expected and taught by Wiggers. *Id.*

Patent Owner proffers several arguments contending that the combination of Amidi and Wiggers fails to teach or suggest a “circuit configured to be mounted on a memory module[] . . . the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal to the common data signal line,” as recited in the challenged claims. PO Resp. 33–35.

Patent Owner first contends that although Wiggers discloses individual switch 29 on a single memory module 24 (as shown in Wiggers Fig. 3), Petitioner identifies only a collective group of switches 29A as its “circuit,” which spans across four (not one) of Wiggers’ memory modules. *Id.* at 33–34. According to Patent Owner, the collective group of switches 29A is designed to function across multiple memory modules and cannot be mounted on only a single memory module. *Id.* at 34. Patent Owner, thus, concludes that the “circuit” in Wiggers fails to meet the claim limitation “circuit configured to be mounted on a memory module,” because the collective group of switches 29A cannot be “configured to be mounted on” a single memory module of Wiggers (memory module 24) or Amidi (e.g., Figures 4A, 4B), even in the Amidi–Wiggers combination. *Id.*

We do not agree with Patent Owner. Rather, we are persuaded by Petitioner’s position and we find that Wiggers’ collective group of switches is controlled by a single memory controller. Ex. 1010, 5:40–51. As can be seen in Figure 5 of Wiggers, there is a single substrate 52 affixed with multiple switches 57 that connect via lines 59 to multiple memory devices

IPR2014-00882
Patent 7,881,150 B2

55. *Id.* at Fig. 5; 6:36–49. The memory module 51 of Figure 5 is a plan view of memory module 24 shown in box 50 of Figure 3. *Id.* at 6:33–36. As shown in Figure 3, the electrical connections include interconnects between the main board portion 27 of the data bus and the module portion 28 on the data bus. *Id.* at 5:1–4; Fig. 3. Therefore, we are satisfied that Wiggers teaches a “circuit configured to be mounted on a memory module[] . . . the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal to the common data signal line,” as recited in the challenged claims.

Patent Owner further contends we erred in construing “circuit configured to be mounted on a memory module,” in an unreasonably broad manner. PO Resp. 34. According to Patent Owner, a person of ordinary skill in the art would understand the claim language of “circuit configured to be mounted on a memory module” as Patent Owner proposed: “an entire circuit configured to be mounted on a single memory module.” *Id.*; *see supra*, Section II.A.3. Based on Patent Owner’s proffered claim construction, Patent Owner argues that the entirety of the Petition’s “circuit” is the collective group (of Wiggers’ switches 29A) itself, which structurally spans multiple memory modules, and thus cannot, and is designed to not, be mounted on only a single memory module. PO Resp. 34–35.

As discussed above, we are not persuaded by Patent Owner’s contentions regarding claims construction. To the contrary, we construe “circuit configured to be mounted on a memory module” to encompass “circuitry configured to be mounted on at least a portion of a memory module.” *See supra*, Section II.A.3. Wiggers’ disclosure of switches (i.e.,

IPR2014-00882
Patent 7,881,150 B2

circuitry) mounted on memory modules controlled by a single memory controller falls within the scope of the term “circuit configured to be mounted on a memory module” as we have construed the term.

Patent Owner also contends that Amidi and Wiggers are improperly combined by Petitioner. PO Resp. 47. Patent Owner explains that the Petitioner’s combination is based on impermissible hindsight and would result in an inoperable system, and that the testimony of Dr. Jagannathan is insufficient to overcome the impediments to combining Amidi and Wiggers. *Id.* at 47–50. According to Patent Owner, Amidi’s CPLD output CS signal and control signals for Wiggers’ switches are not equivalent, based on specific differences in respective timing operation and purpose, and a person of ordinary skill in the art would have understood that the specific implementation details, e.g., timing considerations, are critical in evaluating the operability of using Amidi’s CPLD output CS signal and control signals for Wiggers’ switches equivalently. *Id.* at 48 (citing Ex. 2002 ¶ 113). Patent Owner argues that the timing problems that would arise from an Amidi-Wiggers combination would render such a combination inoperable. *Id.* at 48–49.

Patent Owner relies on the Declaration of Dr. Sechen to support its position regarding the inoperability of an Amidi-Wiggers combination. *See* Ex. 2002 ¶¶ 112–118. Dr. Sechen specifically states that:

due to standardized DDR memory device operation, a DDR read (or write) command’s chip-select signal (e.g., Amidi’s CPLD output chip-select signal) does not coincide with its read (or write) data (e.g., to be gated by Wiggers’ switch control signal). Thus, by failing to properly transfer the target data, the Petition’s proposal—DDR chip-select signals from Amidi’s

IPR2014-00882
Patent 7,881,150 B2

CPLD = Wiggers' switch control signals—would malfunction and be inoperable.

Id. ¶ 113.

Patent Owner then argues that an Amidi-Wiggers combination is also inoperable because a DDR chip-select signal (as used in Amidi) is not designed to be a timing signal, whereas a timing signal is required for Wiggers switch 20A to operate properly. PO Resp. 49. Patent Owner supports its position with the Declaration of Dr. Sechen, who testifies that “[f]undamentally, a DDR chip-select signal in itself lacks timing information to indicate when to properly open and close a switch 29A of Wiggers. Thus, a DDR chip-select signal, as output by Amidi's CPLD, would not be a usable signal at all for controlling Wiggers' switch 29A.” Ex. 2002 ¶ 114.

Again, we do not agree with Patent Owner. Rather, we agree with Petitioner's position and we find that Wiggers teaches timing operations for its switches that would have been applicable to Amidi. Additionally, we are persuaded that the timing requirements for memory devices are dictated by the JEDEC standards, which were known to a person of ordinary skill in the art at the time of the '150 patent. *See e.g.*, Ex. 1008 ¶ 7. We credit the testimony of Dr. Jagannathan, which states that Figure 6 of Wiggers “teaches how to time the coupling and decoupling of a select memory device to the data bus when writing to or reading therefrom.” Ex. 1023 ¶ 51 (citing Ex. 1010, Fig. 6). We also credit the testimony of Dr. Jagannathan regarding the JEDEC21C-4.5.7 standard as it relates to the 168 Pin Registered SDRAM DIMM Family of memory devices. *See e.g.*, Ex. 1007 ¶¶ 43–51. Dr. Jagannathan specifically testifies that

There is a Phase-Locked Loop (PLL) clock input provided to the register and to the memory devices. This is depicted in

IPR2014-00882
Patent 7,881,150 B2

JEDEC21C-4.5.7, Figure as “PCK” input to the register, and as “CK0 _ PLL” in the case of the memory devices. A phase-locked loop device receives an input clock signal and generates another clock signal whose phase matches (within tolerance) the input clock. The details of the PLL are specified in JEDEC21C-4.5.7 at p. 4.5.7-8. Specifically, the PLL clock output is depicted as driving a number of SDRAM devices and registers.

Id. ¶ 45.

The concept of a PLL would generally be well understood by one of ordinary skill in the art. For instance, Jacob teaches that “[t]he function of a PLL or DLL, in general, is to synchronize two periodic signals so that a certain fixed amount of phase-shift or apparent delay exists between them. The two are similar, and the terms are often used interchangeably.”

Id. ¶ 46 (citing Ex. 1018, 11).

Thus, we find that a person of ordinary skill in the art would have known how to address the timing and use of chip-select signals so that the teachings of Wiggers would have been applicable to Amidi.

The test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference (*Keller*, 642 F.2d at 425). Rather, the test is what the combined teachings of those references would have suggested to those of ordinary skill in the art. *Id.* We credit the testimony of Dr. Jagannathan, who states that (i) “when Wiggers refers to a ‘signal’ used in any of the circuits taught therein, one of ordinary skill would understand it is ‘a varying electrical impulse that conveys information from one point to another’” (Ex. 1007 ¶ 93) and (ii) “[o]ne of ordinary skill in the art would have been motivated to apply the switch taught by Wiggers to the memory module of Amidi,

IPR2014-00882
Patent 7,881,150 B2

because the switch of Wiggers increases the speed of data propagation.” Ex. 1007 ¶ 96 (citing Wiggers, 3:39–40, 4:27–37). Thus, we are not persuaded that the control signals from Amidi’s CPLD could not be used as control signals for Wiggers’ switch, or that the teachings of Amidi and Wiggers are not combinable.

Based on the evidence of record, we agree with Petitioner’s position that challenged claims 15–17, 22, 24, 26, and 31–33 would have been obvious over Amidi and Wiggers. First, we are persuaded by Petitioner’s reasoning and the evidentiary record that Amidi teaches a circuit mounted on a memory module that is electrically coupled to a first DDR memory device and a second DDR memory device. We are further persuaded that Amidi teaches a circuit with a logic element, a register, and a PLL. We also are persuaded that the teachings of Amidi could have been implemented using the common data signal line and switch system disclosed in Wiggers so that by (i) selectively electrically coupling a first data signal line to the common data signal line and (ii) selectively electrically coupling a second data signal line to the common data signal line the circuit is responsive to a set of input signals. Additionally, we credit the testimony of Dr. Jagannathan that a person of ordinary skill in the art would have had reason to combine the teachings of Amidi with Wiggers, because both references relate to memory devices, both references discuss the problem of reducing the load seen by the memory controller, and both references describe coupling or isolating memory device loads. *See* Ex. 1007 ¶ 96.

Second, the arguments presented by Patent Owner generally attack the references individually, rather than in combination. PO Resp. 33–35, 47–58. Nonobviousness cannot be established by attacking the references

IPR2014-00882
Patent 7,881,150 B2

individually when a challenge is predicated upon a combination of prior art disclosures. *See Merck*, 800 F.2d at 1097; *cf. Keller*, 642 F.2d at 426 (“[O]ne cannot show non-obviousness by attacking references individually where . . . the rejections are based on combinations of references.”). In attacking the references individually, Patent Owner fails to address Petitioner’s actual challenges and establish an insufficiency in the combined teachings of the references. Patent Owner has not convinced us that Petitioner failed to meet its burden to establish a reasonable likelihood it would prevail in showing obviousness of the challenged claims.

Lastly, we note that the testimony of Patent Owner’s Declarant, Dr. Sechen, is based on the claim constructions proffered by Patent Owner and not on the constructions set forth in the Decision to Institute. *Compare Ex. 2002 ¶¶ 38–72, with Dec. to Inst. 7–12.* Although we have reviewed Dr. Sechen’s testimony in detail, for the reason stated above, we considered as relevant the portions of his analysis regarding the prior art and alleged non-obviousness of the claims and accorded the appropriate weight to that particular testimony.

Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that claims 15, 22, and 31 of the ’150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Amidi and Wiggers.

b. Amidi and Wiggers Teach or Suggest All the Recited Limitations of Dependent Claims 16, 17, 24, 26, 32, and 33

Claims 16, 24, and 32 recite “wherein the two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.” Ex. 1001, 43:3–5, 44:1–3, 44:58–60. Dependent claim 17

IPR2014-00882
Patent 7,881,150 B2

recites that the circuit includes “one or more switches selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the one or more switches operatively coupled to the logic element to receive control signals from the logic element.” Dependent claim 33 recites a similar limitation. *Id.* at 43:6–12. Dependent claim 26 further recites that the claimed circuit is “configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.” *Id.* at 44:8–11.

Petitioner contends the combined disclosures of Amidi and Wiggers, as summarized above, teach or suggest each limitation of dependent claims 16, 17, 24, 26, 32, and 33 of the ’150 patent. Pet. 53–57. Patent Owner does not provide separate contentions regarding additional limitations recited in the dependent claims. *See generally* PO Resp.

After careful consideration of the language recited in claims 16, 17, 24, 26, 32, and 33 of the ’150 patent, the Petition, the Patent Owner Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers, we find that one of ordinary skill in the art would have considered these dependent claims obvious over the combination of Amidi and Wiggers. Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that claims 16, 17, 24, 26, 32, and 33 of the ’150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Amidi and Wiggers.

IPR2014-00882
Patent 7,881,150 B2

III. CONCLUSION

We conclude Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of the '150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combinations of (1) Amidi and Klein and (2) Amidi and Wiggers.

IV. ORDER

For the reasons given, it is

ORDERED that, by a preponderance of the evidence, claims 15–17, 22, 24, 26, and 31–33 of the '150 patent are unpatentable; and

FURTHER ORDERED that because this is a Final Written Decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2014-00882
Patent 7,881,150 B2

FOR PETITIONER:

Fabio Marino
fmarino@mwe.com

Bernie Knight
bknigh@mwe.com

Barrington Dyer
bdyer@mwe.com

MCDERMOTT WILL & EMERY

FOR PATENT OWNER:

Mehran Arjomand
marjomand@mofo.com
patentdocket@mofo.com

David Kim
dkim@mofo.com
patentdocket@mofo.com

Erol Basol
ebasol@mofo.com
patentdocket@mofo.com

Jean Nguyen
jnguyen@mofo.com
patentdocket@mofo.com

Jonathan Statman
jstatman@mofo.com
patentdocket@mofo.com

MORRISON & FOERSTER LLP

Trials@uspto.gov
571.272.7822

Paper No. 34
Filed: December 14, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DIABLO TECHNOLOGIES, INC.,
Petitioner,

v.

NETLIST, INC.,
Patent Owner.

Case IPR2014-01011
Patent 7,881,150 B2

Before LINDA M. GAUDETTE, BRYAN F. MOORE, and
GEORGIANNA W. BRADEN, *Administrative Patent Judges*.

BRADEN, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318 and 37 C.F.R. § 42.73

IPR2014-01011
Patent 7,881,150 B2

I. INTRODUCTION

We have jurisdiction to hear this *inter partes* review under 35 U.S.C. § 6(c). This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of US Patent No. 7,881,150 B2 (Ex. 1001, “the ‘150 patent”) are unpatentable.

A. Procedural History

Diablo Technologies, Inc. (“Petitioner”) filed a Corrected Petition (Paper 5, “Pet.”) to institute an *inter partes* review of claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent. Netlist, Inc. (“Patent Owner”) filed a Preliminary Response (Paper 9, “Prelim. Resp.”). Pursuant to 35 U.S.C. § 314(a), we instituted an *inter partes* review of all challenged claims on the following grounds alleged in the Petition.

References	Basis	Claims Challenged
Ludwig ¹ and Amidi ²	§ 103	15–17, 22, 24, 26, and 31–33
Amidi	§ 102(e)	22, 24, and 26

Paper 12 (“Dec. to Inst.”), 29.

After institution of trial, Patent Owner filed a Patent Owner Response (Paper 26, “PO Resp.”), to which Petitioner filed a Reply (Paper 28, “Reply”). An oral argument was held on July 28, 2015, consolidated with

¹ US Patent No. 5,581,498, iss. Dec. 3, 1996 (filed Oct. 20, 1994) (“Ludwig,” Ex. 1011).

² US Patent Publication No. 2006/0117152 A1, pub. June 1, 2006 (filed Jan. 5, 2004) (“Amidi,” Ex. 1008).

IPR2014-01011
Patent 7,881,150 B2

the oral hearings for IPR 2014-00882 and IPR2014-00883. *See* Paper 31. A transcript (“Tr.”) of the oral argument is included in the record. Paper 32.

B. Related Proceedings

The parties informs us that the ’150 patent is involved in the following federal district court cases: *Diablo Technologies, Inc. v. Netlist, Inc.*, Case No. 4:13-CV-03901-YGR (N.D. Cal.); and *Netlist, Inc. v. Smart Modular Technologies*, Case No. 4:13-CV- 05889-YGR (N.D. Cal.). Papers 10, 1; 33, 2. In addition, Petitioner filed two other petitions requesting *inter partes* review of the ’150 patent. Paper 10, 2. These cases are: IPR 2014-00882 and IPR2014-00883. *Id.* We consolidated the oral hearings for IPR2014-00882, IPR2014-00883, and IPR 2014-01011. *See* Paper 31.

Petitioner further informs us that related US Patent Nos. 7,619,912 and 7,636,274 are the subjects of *inter partes* reexaminations (95/000,578 and 95/001,337). Pet. 10–11. Petitioner also informs us that related US Patent No. 7,289,386 is the subject of district court case *Google, Inc. v. Netlist, Inc.*, Case No. C 08-4144-SBA (N.D. Cal.). *Id.* at 14–15.

C. The ’150 Patent

The ’150 patent relates to a memory module of a computer system with improved performance and memory capacity. Ex. 1001, 1:30–34. Memory module 10 includes a plurality of memory devices 30 (arranged in ranks 32) and circuit 40. *Id.* at 4:56–65; Fig. 1. Circuit 40 is electrically coupled to the memory devices 30 and memory controller 20 of a computer system. *Id.* The memory module improves performance and memory capacity by isolating electrical loads of the memory devices from the computer system. *Id.* at 4:65–66.

IPR2014-01011
Patent 7,881,150 B2

Circuit 40 receives input signals from memory controller 20. *Id.*

Figure 1, reproduced below, illustrates input signals (corresponding to a number of memory devices) from memory controller 20, such as chip select signals (“cs#”), that are directed to memory module 10, which can act as a virtual memory module. *Id.* at 16:47–57; Figs. 1, 9A, 9B.

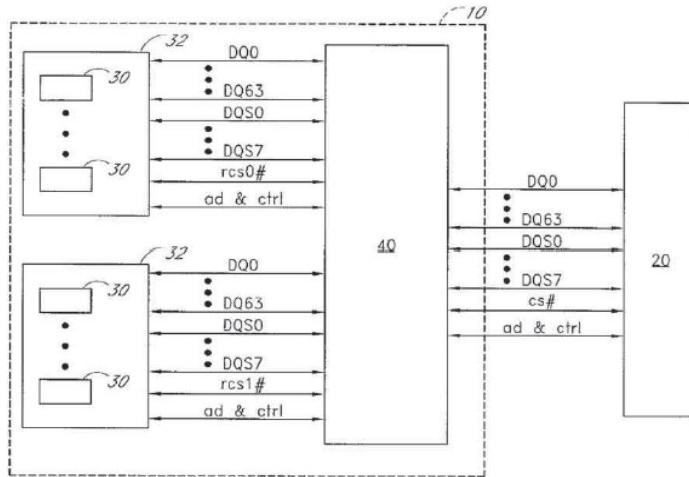


Figure 1 is a schematic of a memory module with circuit 40 and memory devices 30 and connectable to memory controller 20.

As shown in Figure 1 above, based on the received input signals from memory controller 20, circuit 40 generates output signals corresponding to memory devices 30 on the memory module. *Id.* at Fig. 1. The output signals include a different number of chip select signals (e.g., “rcs0#” and “rcs1#”) corresponding to memory devices 30 shown in ranks 32. *Id.* at 16:66–17:4; Fig. 1.

Circuit 40 includes a logic element, such as a programmable logic device (PLD), an application-specific integrated circuit (ASIC), a field programmable gate array (FPGA), or a complex programmable-logic device (CPLD). Ex. 1001, 6:4–18. As shown in Figure 9A, reproduced below, circuit 40 may also include register 230 and phase-lock loop device (PLL) 220. *Id.* at 15:35–41; Fig. 9A.

IPR2014-01011
Patent 7,881,150 B2

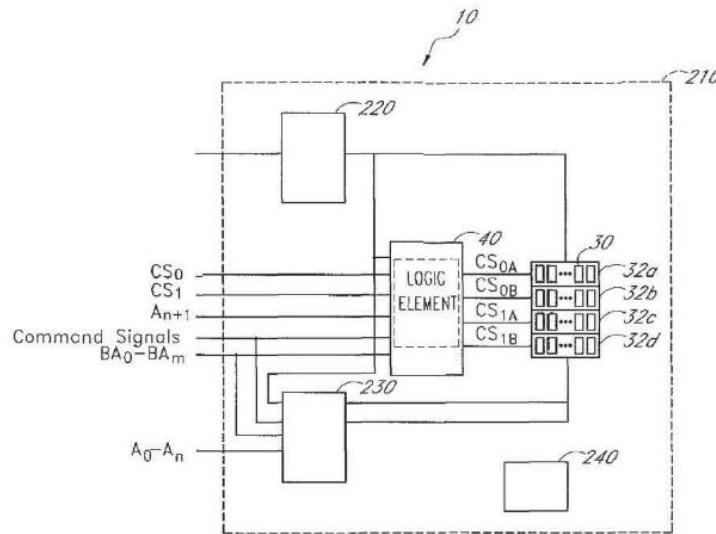


Figure 9A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Figure 9A above illustrates circuit 40 receiving a set of input command signals, address signals (A_{n+1}), including bank address signals (BA_0-BA_m), row address signals (A_0-A_n), column address signals, gated column address strobe signals, and chip-select signals (CS_0, CS_1), from memory controller 20 of the computer system. *Id.* at 16:24–29, 17:11–26. “In response to the set of input address and command signals, circuit 40 generates a set of output address and command signals.” *Id.* at 16:31–33.

With the output address and command signals, circuit 40 isolates the electrical loads of some memory devices 30 from the computer system. *Id.* at 6:48–62. According to the ’150 patent, load isolation may result in specific benefits including reduced load related to data signal lines. *Id.* at 14:34–40. In order to isolate the loads, the logic element of circuit 40 translates between a system memory domain of the computer system and a physical memory domain of memory module 10. *Id.* at 6:48–62. As shown in Figure 3, reproduced below, the circuit isolates the load of a memory

IPR2014-01011
Patent 7,881,150 B2

device by isolating one or both of DQ data signal lines 102a, 102b of two memory devices 30a and 30b from common DQ data signal line 112 that is coupled to the computer system. *Id.* at 6:63–7:2, Fig. 3A.

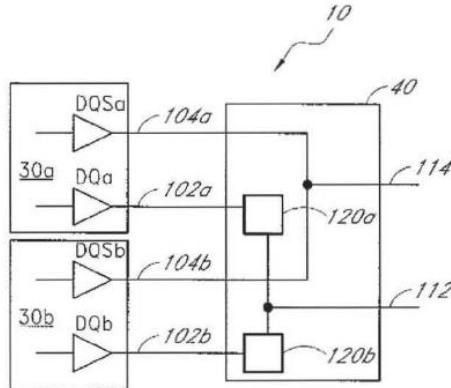


Figure 3A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Circuit 40, shown in Figure 3A above, can electrically couple one or both of DQ data signal lines 102a, 102b of memory devices 30a and 30b to common data signal line 112, at the same time. *Id.* at 7:22–26; Fig. 3A.

Circuit 40 also allows a DQ data signal to be transmitted from memory controller 20 of the computer system to one or both of DQ data signal lines 102a, 102b. *Id.* at 7:2–5. The logic element of circuit 40 uses switches 120a, 120b in order to isolate or couple one or both of DQ data signal lines 102a, 102b of memory devices 30a and 30b from common data signal line 112. *Id.* at 7:2–12.

D. Illustrative Claim

As noted above, *inter partes* review was instituted for claims 15–17, 22, 24, 26, and 31–33 of the '150 patent, of which claims 15, 22, and 31 are independent claims. Claim 15 is illustrative of the challenged claims and is reproduced below:

IPR2014-01011
Patent 7,881,150 B2

15. A circuit configured to be mounted on a memory module so as to be electrically coupled to a first double-data-rate (DDR) memory device having a first data signal line and a first data strobe line, to a second DDR memory device having a second data signal line and a second data strobe line, and to a common data signal line, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;

a register;

a phase-lock loop device configured to be operationally coupled to the first DDR memory device, the second DDR memory device, the logic element, and the register,

wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the memory module, wherein the system memory domain has a first memory density per memory device, and the physical memory domain has a second memory density per memory device less than the first memory density per memory device.

Ex. 1001, 42:41–43:2.

II. DISCUSSION

A. *Claim Construction*

In an *inter partes* review, claim terms in an unexpired patent are interpreted according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b); *see also In re Cuozzo Speed Techs., LLC*, 793 F.3d 1268, 1278–79 (“Congress implicitly approved the broadest reasonable interpretation standard in

IPR2014-01011
Patent 7,881,150 B2

enacting the AIA,” and “the standard was properly adopted by PTO regulation.”). Under this standard, claim terms generally are given their ordinary and customary meaning, as understood by one of ordinary skill in the art in the context of the patent’s entire written disclosure. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). Yet a “claim term will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.” *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002).

In the Decision to Institute, we construed the terms “Memory Module,” “Circuit Configured to be Mounted on a Memory Module,” and “Selectively Electrically Coupling,” which are recited in all the challenged independent claims. *See* Dec. to Inst. 8–12. During the course of the trial, Patent Owner argued for altered constructions of these claim terms. PO Resp. 4–15. Therefore, we address these contentions and construe each claim term as discussed below.

1. “Memory Module”

In the Decision to Institute, we construed the term “memory module,” as “a plurality of memory devices and a circuit” thereby encompassing “additional circuitry and multiple printed circuit boards.” Dec. to Inst. 8–9.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 2; Tr. 5:25–6:10. Patent Owner, however, contends that “memory module” should be construed as “a packaging arrangement of one or more memory device(s) for use in a computer socket.” PO Resp. 4; Tr. 47:5–7. According to Patent Owner, the construction of “memory module” in the Decision to Institute is unreasonably broad, because it is inconsistent

IPR2014-01011
Patent 7,881,150 B2

with the ordinary and customary meaning of the term as it would be understood by a person of ordinary skill in the art. PO Resp. 4–10. Patent Owner argues the Board erred in construing the term by analyzing each component of the word separately (*id.* at 5) and relying on the '150 patent specification (Tr. 47:17–20), whereas a person of ordinary skill in the art would have understood “memory module” to be a term of art (PO Resp. 5–6; Ex. 2002 ¶ 52 (Declaration of Dr. Carl Sechen)). Patent Owner explains that under the Board’s construction of “memory module,” the term would encompass a memory controller and associated memory devices. PO Resp. 9–10.

Patent Owner further contends that the Board’s construction of “memory module” is inconsistent with the '150 patent disclosure. *Id.* at 10. Patent Owner notes that claims 15, 22, and 31 recite “the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals.” *Id.* (emphasis omitted). According to Patent Owner, due to the use of different terms in the '150 patent, “memory module” would not be read as including the '150 patent’s “memory controller” by a person of ordinary skill in the art. *Id.*; Ex. 2002 ¶ 57.

We are charged with interpreting claim terms according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b). Additionally, when construing claim terms, we “should also consult the patent’s prosecution history in proceedings in which the patent has been brought back to the [U.S. Patent and Trademark Office] for a second review.” *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015). Yet, we must be careful not to

IPR2014-01011
Patent 7,881,150 B2

improperly import limitations into the claims or to read a particular embodiment appearing in the written description into the claim, if the claim language is broader than the embodiment. *In re Van Geuns*, 988 F.2d 1181, 1184 (Fed. Cir. 1993).

The specification of the '150 patent does not define explicitly the term "memory module." The specification does, however, teach embodiments that describe a memory module as comprising a plurality of memory devices on a carrier and a circuit. Ex. 1001, 2:63–64; 3:7–9; 4:59–63. In another embodiment, a memory module comprises (i) a printed circuit board on which memory devices are mounted, (ii) a plurality of edge connectors configured to be electrically coupled to a corresponding plurality of contacts of a module slot of the computer system, and (iii) a plurality of electrical conduits which electrically couple the memory devices to the circuit and which electrically couple the circuit to the edge connectors. *Id.* at 5:24–32. The '150 patent also teaches that memory modules in the disclosed embodiments are compatible with at least single in-line memory modules (SIMMS) and dual in-line memory modules (DIMMS). *Id.* at 5:32–39.

Although the embodiments disclosed in the '150 patent are instructive, the claims recite language broader than that found in the embodiments. *See In re Van Geuns*, 988 F.2d at 1184. Therefore, we decline to adopt Patent Owner's claim construction as it would import limitations improperly from the specification into the claims and unnecessarily limit the scope of the claims. We credit, however, the testimony of Patent Owner's Declarant, Carl Sechen, Ph.D. ("Dr. Sechen"), who explains the state of the art and the customary meaning of "memory module" as it would be understood by one of ordinary skill in the art to

IPR2014-01011
Patent 7,881,150 B2

encompass at least a “removable circuit board, cartridge, or other carrier that contains one or more RAM memory chips.” *See Ex. 2002 ¶¶ 39–57.*

Therefore, we modify the construction of “memory module” from that set forth in the Decision to Institute, wherein we construed the term as “a plurality of memory devices and a circuit” that “encompass[es] additional circuitry and multiple printed circuit boards.” Dec. to Inst. 9. Rather, we construe the term “memory module” as “one or more memory devices on a carrier,” because such a construction is consistent with the disclosure of the ’150 patent and with the ordinary and customary meaning of “memory module.”

2. *“Selectively Electrically Coupling” and “Selectively Electrically Isolating”*

In the Decision to Institute, we construed the term “selectively electrically coupling,” as “making a selection between at least two components so as to transfer power or signal information from one component to at least one other component.” Dec. to Inst. 9–11.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 2; Tr. 5:25–6:10. Patent Owner, however, contends that “selectively electrically coupling” should be construed as “electrically coupling in response to a selection.” PO Resp. 12–15; Tr. 70:4–9. According to Patent Owner, the Board’s construction is unreasonably broad, whereas its proffered construction is more consistent with the disclosure of the ’150 patent. PO Resp. 13–14 (citing Ex. 1001, 5:29–30, 7:22; Ex. 2002 ¶ 66), 26–28. Patent Owner specifically argues that “electrically coupling” in the ’150 Patent is provided by a structural pathway for electricity, and this is also consistent with the meaning of “electrically coupling” as a term of art. *Id.* at 14 (citing Ex. 2002 ¶ 66). Patent Owner further argues that a

IPR2014-01011
Patent 7,881,150 B2

person of ordinary skill in the art would understand the act of electrically coupling to take place between strictly two components, between which a structural pathway for electricity would be formed. *Id.*

We are unpersuaded by Patent Owner's position. The specification of the '150 patent does not define explicitly the term "selectively electrically coupling." Therefore, we refer to its ordinary and customary meaning, as would be understood by one of ordinary skill in the art in the context of the entire disclosure. *In re Translogic Tech., Inc.*, 504 F.3d at 1257. A technical dictionary, IEEE Dictionary³, defines "electrical coupling" as "[e]lectrical charges in conductors of a disturbed circuit formed by electrical induction." Ex. 3001. The IEEE Dictionary explains that "[s]ince the ratio of a conductor's electrostatic charge to the potential difference between conductors (required to maintain that charge) is the general definition of capacitance, electrical coupling is also called capacitive coupling." *Id.* The IEEE Dictionary defines "coupling capacitance (1) (ground systems)" ("capacitive coupling") as "[t]he association of two or more circuits with one another by means of capacitance mutual to the circuits." Ex. 3002. We understand this to mean that the two or more circuits are associated in such a way that power or signal information may be transferred from one circuit to another. The Oxford English Dictionary defines "selectively" as "[i]n a selective manner; by selection." Ex. 3003. The Oxford English Dictionary also defines "select" as "[t]o choose or pick out in preference to another or others." Ex. 3004.

³ IEEE 100: The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition, Standards Information Network, IEEE Press (2000).

IPR2014-01011
Patent 7,881,150 B2

Accordingly, we modify slightly the construction from the Decision to Institute of “selectively electrically coupling,” as “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component,” because such a construction is consistent with the ordinary and customary meaning of “selectively electrically coupling.”

Based on the same reasoning, we construe “selectively electrically isolating” as “making a selection between at least two components and not transferring power or signal information from one selected component to the other selected component.”

3. “Circuit Configured to be Mounted on a Memory Module”

In the Decision to Institute, we construed the term “a circuit configured to be mounted on a memory module,” to encompass “circuitry configured to be mounted on at least a portion of a memory module.” Dec. to Inst. 11–12. We determined such construction is consistent with the ordinary and customary meaning of “a circuit configured to be mounted on a memory module.” *Id.* at 12.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 2; Tr. 5:25–6:10. Patent Owner, however, contends that “a circuit configured to be mounted on a memory module” should be construed as “an entire circuit configured to be mounted on a single memory module.” PO Resp. 18–19; Tr. 68:1–18.

Patent Owner notes that our claim construction, as set forth in the Decision to Institute, is ambiguous in that it can be read two ways:

One might read the Board’s construction as meaning that the circuit is mounted on and occupies at least a portion of the memory module (Ex. 2002, ¶ 62), which may be consistent with

IPR2014-01011
Patent 7,881,150 B2

Netlist's construction. On the other hand, one might read the Board's construction as encompassing portions of the circuit to be mounted off-module, which would be unreasonably broad to a [person of ordinary skill in the art].

Id. at 11 (citing Ex. 2002 ¶ 62). According to Patent Owner, “memory module” is a term of art that would have had a well-understood meaning to a person of ordinary skill in the art at the time of the invention and a person of ordinary skill in the art would not have understood “a circuit configured to be mounted on a memory module” to include circuit parts off of that memory module or on a different memory module. *Id.* at 12 (citing Ex. 2002 ¶ 60). Patent Owner contends that such a construction could include situations that defeat the purpose of a memory module to make removing and installing memory upgrades easy and error-free. *Id.* (citing Ex. 2002 ¶ 61).

We decline to adopt Patent Owner's claim construction as it is inconsistent with the definition of “circuit” as found in the specification of the '150 patent. The '150 patent defines “circuit” as “a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.” Ex. 1001, 5:9–13. The '150 patent does not limit a “circuit” to only a configuration of electrical components or devices that are mounted on a single memory module. Therefore, applying the broadest reasonable interpretation consistent with the specification of the '150 patent, we construe the claim element “a circuit configured to be mounted on a memory module,” as we did in the Decision to Institute, but we further

IPR2014-01011
Patent 7,881,150 B2

clarify the construction to encompass “a portion of circuitry configured to be mounted on at least a portion of a memory module.”

4. Other Claim Terms

We determine that no express constructions of any other claims terms are required for our analysis, and we apply the ordinary and customary meaning of each claim term.

B. Principles of Law

To prevail in its challenges to the patentability of the claims, a petitioner must establish facts supporting its challenges by a preponderance of the evidence. 35 U.S.C. § 316(e); 37 C.F.R. § 42.1(d).

A claim is unpatentable under 35 U.S.C. § 102 if a prior art reference discloses every limitation of the claimed invention, either explicitly or inherently. *Glaxo Inc. v. Novopharm Ltd.*, 52 F.3d 1043, 1047 (Fed. Cir. 1995). Furthermore, the prior art reference—in order to be anticipatory—must disclose every limitation of the claimed invention arranged or combined in the same way, as in the claim. *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1371–72 (Fed. Cir. 2008). A reference can anticipate a claim, however, even if it “‘d[oes] not expressly spell out’ all the limitations arranged or combined as in the claim, if a person of skill in the art, reading the reference, would ‘at once envisage’ the claimed arrangement or combination.” *Kennametal, Inc. v. Ingersoll Cutting Tool Co.*, 780 F.3d 1376, 1381 (Fed. Cir. 2015) (quoting *In re Petering*, 49 CCPA 993, 301 F.2d 676, 681 (1962)). Additionally, “the reference need not satisfy an *ipsissimis verbis* test.” *In re Gleave*, 560 F.3d 1331, 1334 (Fed. Cir. 2009) (internal citation omitted).

IPR2014-01011
Patent 7,881,150 B2

A claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

We analyze the instituted ground of unpatentability in accordance with the above-stated principles.

C. Level of Ordinary Skill in the Art

In determining whether an invention would have been obvious at the time it was made, we determine the level of ordinary skill in the pertinent art at the time of the invention. *Graham v. John Deere*, 383 U.S. at 17. “The importance of resolving the level of ordinary skill in the art lies in the necessity of maintaining objectivity in the obviousness inquiry.” *Ryko Mfg. Co. v. Nu-Star, Inc.*, 950 F.2d 714, 718 (Fed. Cir. 1991).

Petitioner’s Declarant, Srinivasan Jagannathan, Ph.D. (“Dr. Jagannathan”), testifies that a person of ordinary skill in the art at the time of the ’150 patent:

would understand basic memory and data communication concepts, with a bachelor’s degree in any of electrical engineering, computer engineering, computer science, or related field. Course work for one of ordinary skill in the art

IPR2014-01011
Patent 7,881,150 B2

would have included a course on computer organization, principles of digital design, or computer architecture. One of ordinary skill in the art would also have around one year of experience related to computer memory systems. For example, such experience may include experience in DRAM memory technology and related industry standards such as JEDEC standards for DRAM memories and memory modules.

Ex. 1007 ¶ 53. Patent Owner' Declarant, Dr. Sechen, testifies that one of ordinary skill in the art at the time of the '150 patent would have had an undergraduate degree in electrical engineering or computer engineering, at least two years of professional experience in the design of memory systems, familiarity with the latest JEDEC standard specifications for memory devices and modules, and familiarity with the latest DRAM memory devices widely available in the market. Ex. 2002 ¶ 14. Dr. Sechen further testifies that a person of ordinary skill in the art would have design proficiency in memory modules comprising double-data rate (DDR) memory technology, such as memory modules with JEDEC standard DDR SDRAM devices. *Id.* ¶ 15.

Based on our review of the '150 patent and the types of problems and solutions described in the '150 patent and cited prior art, we conclude a person of ordinary skill in the art at the time of the '150 patent would have a Bachelor's degree in electrical engineering, computer engineering, computer science, or related field, and at least one year of work experience, including familiarity with computer memory systems and related industry standards such as JEDEC standards for DRAM memories and memory modules. We further note that the applied prior art reflects the appropriate level of skill at

IPR2014-01011
Patent 7,881,150 B2

the time of the claimed invention. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001).

D. Expert Testimony

Patent Owner argues that Petitioner's Declarant, Dr. Jagannathan, does not qualify as a person of ordinary skill in the art at the time of the invention because of his alleged lack of experience designing memory modules. PO. Resp. 16–17. According to Patent Owner, Dr. Jagannathan's experience and background is directed to software and is not relevant to the case. *Id.* at 18–23. Patent Owner argues that, unlike Dr. Jagannathan, its Declarant Dr. Sechen has significant practical experience designing memory modules. *Id.* at 18; Ex. 2001 ¶¶ 3, 4, Exhibit A.

As to his hardware and memory design experience, Dr. Jagannathan was awarded a Doctorate degree in Computer Science, and has “two decades of experience in the design, development, and analysis of a wide range of hardware, software, network and database systems.” Ex. 1007 ¶ 2. He has designed and implemented hardware virtual memory caches, and researched theoretical performance measures of various cache coherency protocols. *Id.* Dr. Jagannathan also stated during his deposition that he has experience in the design of memory systems. Ex. 2003, 124:12–126:1. Patent Owner contends, however, that Dr. Jagannathan fails to qualify as a person of ordinary skill in the art because he lacks sufficient professional experience physically designing (i.e., “actually putting down a design and saying this is what it would be”) a memory module. PO Resp. 21 (citing Ex. 2003, 125:14–17). We disagree.

To testify as an expert under FRE 702, a person need not be a person of ordinary skill in the art, but rather must be “qualified in the pertinent art.”

IPR2014-01011
Patent 7,881,150 B2

Sundance, Inc. v. DeMonte Fabricating Ltd., 550 F.3d 1356, 1363–64 (Fed. Cir. 2008); *see SEB S.A. v. Montgomery Ward & Co.*, 594 F.3d 1360, 1372–73 (Fed. Cir. 2010) (upholding a district court’s ruling to allow an expert to provide testimony at trial because the expert “had sufficient relevant technical expertise” and the expert’s “knowledge, skill, experience, training [and] education . . . [wa]s likely to assist the trier of fact to understand the evidence”); *Mytee Prods., Inc. v. Harris Research, Inc.*, 439 Fed. App’x 882, 886–87 (Fed. Cir. 2011) (non-precedential) (upholding admission of the testimony of an expert who “had experience relevant to the field of the invention,” despite admission that he was not a person of ordinary skill in the art). We find that, although Dr. Jagannathan is less experienced than Dr. Sechen in the area of memory module design, he is qualified sufficiently to testify about memory systems and memory modules.

E. Alleged Obviousness of Claims 15–17, 22, 24, 26, and 31–33 in view of Ludwig and Amidi

Petitioner alleges claims 15–17, 22, 24, 26, and 31–33 of the ’150 patent are unpatentable under 35 U.S.C. § 103 over the combination of Ludwig and Amidi. Pet. 21–40. Patent Owner disputes Petitioner’s position, arguing that a person of ordinary skill in the art would not have had reason to combine the references in the manner proposed by Petitioner (PO Resp. 39–44) and further that the combination of the references fails to teach or suggest all of the claim limitations (*id.* at 24–26, 32–39).

We have reviewed the Petition, the Patent Owner Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 15–17, 22, 24, 26, and 31–33 of

IPR2014-01011
Patent 7,881,150 B2

the '150 patent would have been obvious to a person of ordinary skill in the art at the time of the invention in view of the combination of Ludwig and Amidi.

1. Overview of Ludwig

Ludwig discloses the use of integrated circuit (IC) chips stacked together, and which act as a single memory chip (VIC). Ex. 1011, 1:9–12; 1:43–48. Ludwig further discloses that the stacked chips have four memory layers, with each layer including four DRAM memory chips that receive and transmit data using data lines DQ1-DQ4. *Id.* at 9:31–39. One embodiment of the stacked chips is shown in Figure 3, reproduced below.

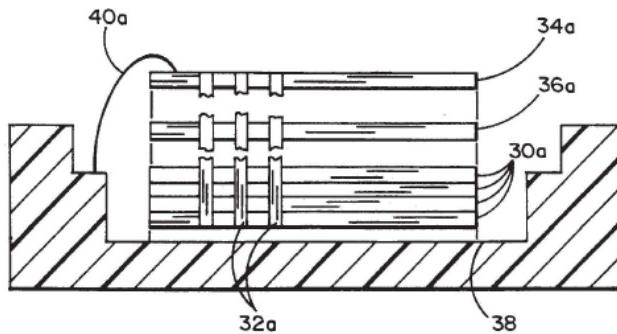


Figure 3 illustrates stacked memory chips 30a with ceramic cap layer 34a and VIC chip 36a located between the four memory chips.

As shown in Figure 3, VIC chip 36a is an interface chip that is electronically interposed between memory chips 30a and the host system and provides appropriate connections between memory chips 30a and the host system. *Id.* at 2:6–7; 4:60–65; 5:9–11; 5:36–43; 5:64–67.

Another embodiment of the stacked chips is shown in Figure 7, reproduced below.

IPR2014-01011
Patent 7,881,150 B2

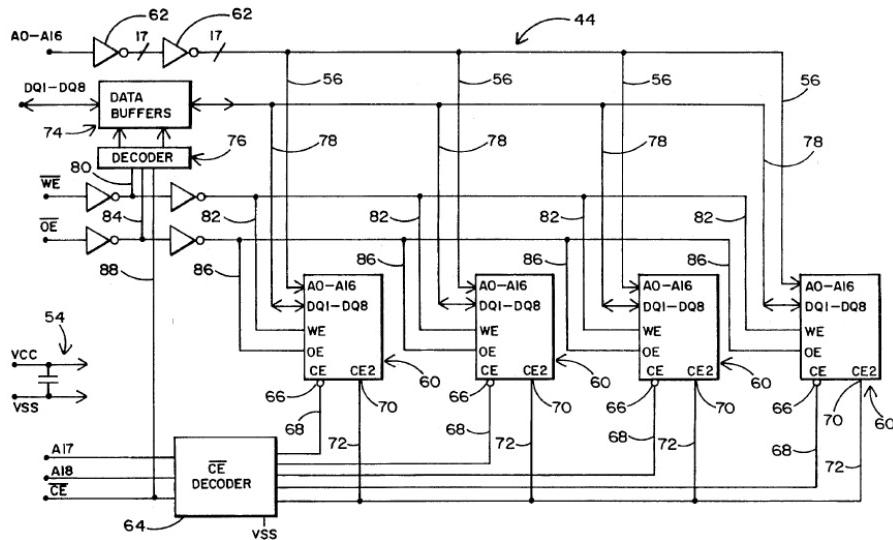


Figure 7 illustrates the circuitry associated with the stacked chips, including the four IC memory chips and VIC chip 44. *Id.* at 5:63–65.

VIC chip 44, shown in Figure 7, provides connections between the host system and stacked memory chips 60 (*id.* at 5:65–67), while the four IC memory chips provide memory capacity (*id.* at 6:11–12). According to Ludwig, “recapitulating the use of the VIC chip to cause the four chip stack to be addressed as if there were a single higher capacity chip, there are address lines (A0-A18) available at the chip stack.” *Id.* at 6:28–36. Ludwig discloses that A17 and A18 are decoded in the VIC chip to select one of the four chips in the stack, while the remaining address lines, A0-A16, feed into all four memory chips. *Id.* at 6:31–33. Thus, per this embodiment in Ludwig, A17 and A18 select the stack layer and A0-A16 select the memory location. *Id.* at 6:34–35.

Ludwig further discloses that another portion of the memory interface is focused on data transmission, which can flow in both directions, i.e., from the host “system into the stacked chip package, or out of the package into the host system.” *Id.* at 6:40–43. As shown in Figure 7 above, data buffer decoder 76 controls data flow, while the data travels along 8 parallel lines,

IPR2014-01011
Patent 7,881,150 B2

DQ1—DQ8. *Id.* at 6:44–51. According to Ludwig, these data lines in the VIC chip are buffered by “tri-state” buffers 74. *Id.*

Another embodiment of the stacked chips is shown in Figure 11, reproduced below.

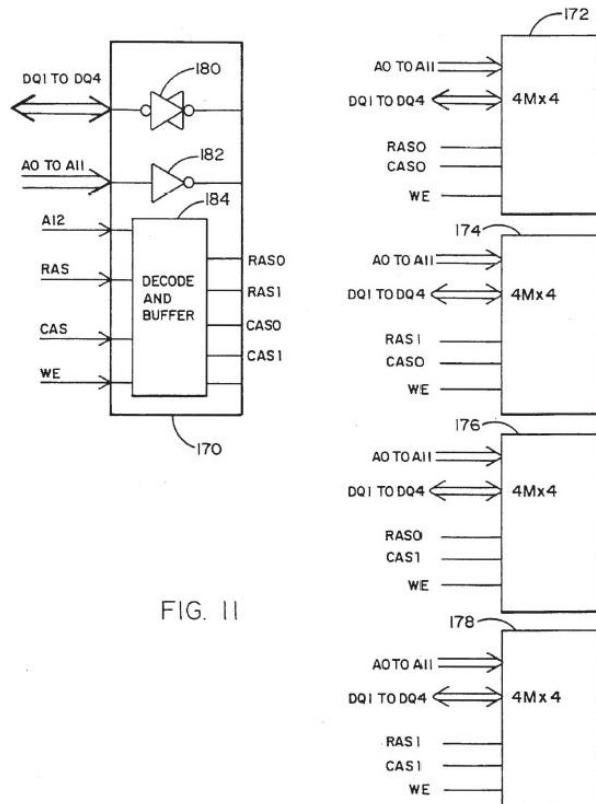


FIG. 11

Figure 11 illustrates a diagram of DRAM memory package organized with a VIC chip and four stack DRAM IC chips.

As illustrated in Figure 11, stacked memory chips 172, 174, 176, 178 are connected to data lines DQ1–DQ4. *Id.* at 9:37–39. VIC chip 170 provides address decoding, so that decoding of data takes place by routing the input of VIC buffers to the appropriate memory chip in the stack. *Id.* at 9:13–23.

2. Overview of Amidi

Amidi discloses a memory interface system with a processor, a memory controller, and a memory module. Ex.1008 ¶¶ 2, 3. According to

IPR2014-01011
Patent 7,881,150 B2

Amidi, a prior art memory interface system is shown in Figure 1, reproduced below.

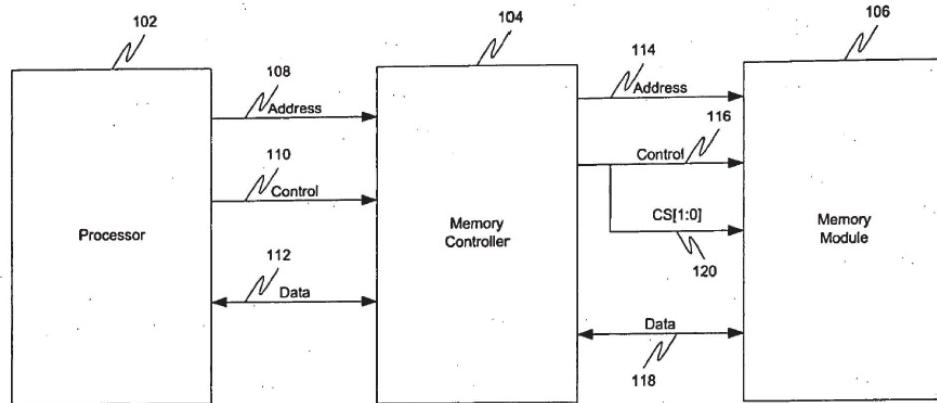


Figure 1 is a schematic of a standard prior art memory interface system.

The prior art system in Figure 1 includes memory module 106 with controller address bus 114, controller control signal bus 116, and controller data bus 118. *Id.* ¶ 2, Fig. 1. As illustrated in Figure 1, memory module 106 communicates with memory controller 104 via busses 114, 116, 118. *Id.* at Fig. 1. Amidi teaches that each stack of DDR memory devices has a data signal line and a data strobe line DQS. *Id.* ¶ 32; Fig. 2. Amidi also teaches that at least two DDR memory devices are connected to a common data memory bus. *Id.* ¶ 34; Fig. 3.

Amidi further discloses multiple memory devices mounted on the front and back side of memory module 400 as shown in Figure 4A reproduced below. *Id.* ¶¶ 34, 37.

IPR2014-01011
Patent 7,881,150 B2

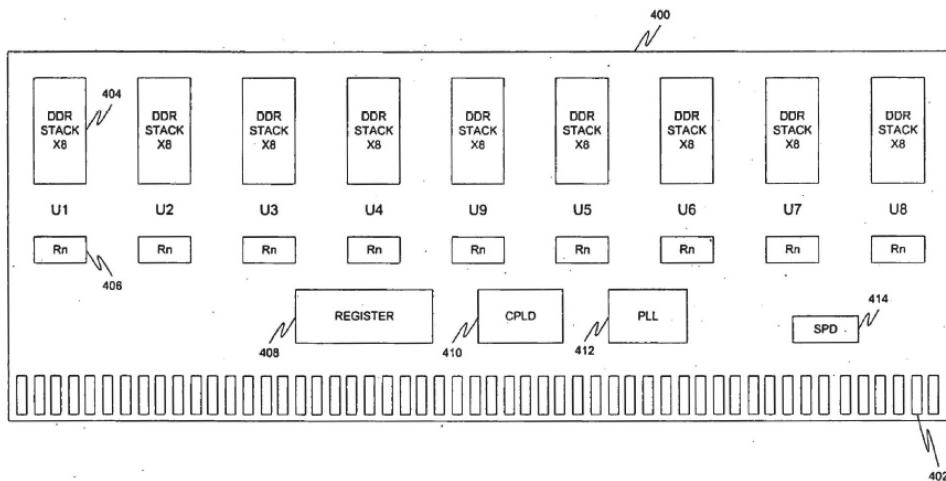


Figure 4A is a schematic of a DDR memory module.

Figure 4A, above, illustrates one embodiment of Amidi where memory module 400 includes memory devices 404, resistor network 406, register 408, complex programmable logic device (CPLD) 410, phase-locked loop (PLL) 412, and SPD 414⁴. *Id.* According to Amidi, memory module 400 receives input signals, including address (Add(n)) signals, row address strobe (RAS) signal, column address strobe (CAS) signal, and bank address (BA[1:0]) signals. Ex. 1008 ¶ 50; Fig. 6A.

Another embodiment of Amidi's memory interface system is shown in Figure 6A, reproduced below.

⁴ According to Amidi, SPD 414 is a simple “interface Electrically Erasable Programmable Read-Only Memory (EEPROM) to hold information regarding memory module for BIOS during the power-up sequence.” Ex. 1008 ¶ 40.

IPR2014-01011
Patent 7,881,150 B2

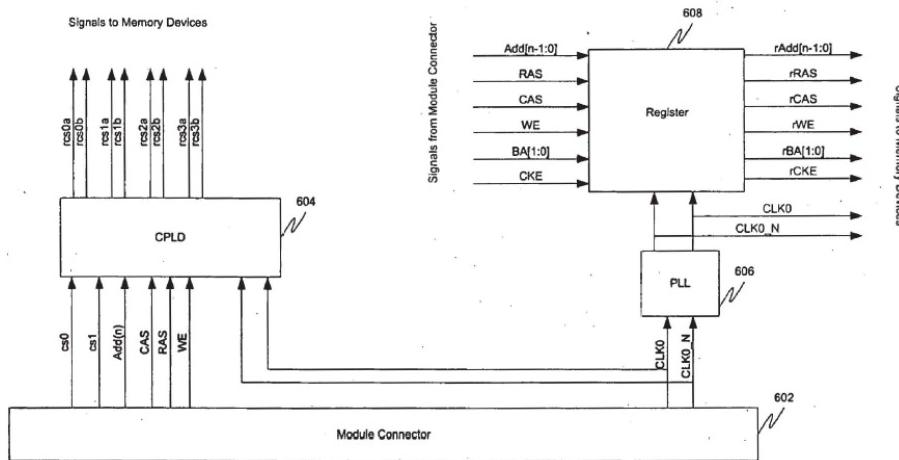


Figure 6A is a schematic of a row address decoding system for a transparent four rank memory module.

As illustrated in Figure 6A above, module connector 602 sends signals to CPLD 604, PLL 606, and register 608. *Id.* CPLD 604 also ensures that all commands for a two rank memory module conveyed by module connector 602 are performed on the four rank memory modules. *Id.* ¶ 52. Amidi explains that the system chip select signals control the ranks of individual memory modules. *Id.* ¶¶ 2, 3.

3. Analysis

a. Ludwig and Amidi Teach or Suggest All the Recited Limitations of Independent Claims 15, 22, and 31

Petitioner contends the combined disclosures of Ludwig and Amidi, as summarized above, teach or suggest each limitation of independent claims 15, 22, and 31 of the '150 patent. Pet. 21–40. Petitioner first argues that Ludwig discloses an interface chip that qualifies as a circuit that is mounted on a stacked chip memory module, where the stack is a plurality of DDR memory devices arranged in one or more ranks. Pet. 21–23 (citing Ex. 1011, 2:1–9; 12:29–44). According to Petitioner, the VIC chip (“circuit”) of Ludwig includes logic elements, a register, a PLL, and a memory element in

IPR2014-01011
Patent 7,881,150 B2

the form of a memory cache. *Id.* at 21–22 (citing Ex. 1011, 7:1–4; 7:16–19; 7:44–55; Figs. 7, 8); Ex. 1007 ¶¶ 103, 106, 107, 108. Petitioner then explains that Amidi discloses a circuit mounted on a memory module, where the circuit includes a logic element, a register, and a phase-lock loop device. *Id.* at 22 (citing Ex. 1008, ¶¶ 2, 37; Figs. 4A, 6A); Ex. 1007 ¶¶ 58, 64. Petitioner also contends that the system described in Amidi includes a memory module having one or more ranks of DDR memory devices, which are electrically coupled to the components of the circuit (CPLD). *Id.* at 22–23 (citing Ex. 1008, Figs. 4A, 4B, 6A.)

Petitioner then argues that Ludwig discloses memory devices having data signal lines and data strobe lines. Petitioner specifically points out that Ludwig discloses a strobe signal used in the form of combinations of CAS with WE and OE pins. *Id.* at 25 (citing Ex. 1011, 8:44–56; 6:50–52; 7:43–61; Fig. 9). Petitioner notes that Amidi also discloses memory devices having data signal lines and data strobe lines. *Id.* at 27 (citing Ex. 1008 ¶¶ 29, 32; Figs. 2, 3). According to Petitioner, Amidi discloses that each stack of DDR memory devices has a data signal line and a data strobe line DQS. *Id.* at 26 (citing Ex. 1008 ¶ 32; Fig. 3). Amidi also discloses that at least two DDR memory devices are connected to the same (common) memory bus (“common data signal line”). *Id.* (citing Ex. 1008 ¶¶ 34–35; Fig. 3). According to Petitioner, a person of ordinary skill in the art would recognize that each DDR memory device has its own data pins (data bus), that the memory devices are connected to a common data signal line, and that the circuit of Amidi is “electrically coupled” to the common data bus. Pet. 26; *see* Ex. 1007 ¶¶ 63, 72.

IPR2014-01011
Patent 7,881,150 B2

Petitioner provides arguments that the sending and receiving of input signals by the systems occurs in Ludwig and Amidi, and explains how each reference teaches a circuit that is responsive to such input signals. Pet. 27–38. Petitioner specifically argues that Ludwig teaches a circuit that is “responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line.” Petitioner makes this argument because Ludwig discloses decoder 64 is part of the circuit and responds to (i) address lines A17, A18 and chip enable CE as input signals, and (ii) chip select signals CEM 1–4 as output signals to drives buffers 113, 114 to selectively couple data signal line DQ1A–DQ8A of memory chip layer A to common data signal line DQ1–DQ8. *Id.* at 29–30 (citing Ex. 1011, 7:16–48; Fig. 8); Ex. 1007 ¶¶ 105, 112.

Petitioner also argues that Amidi teaches “selectively coupling” because Amidi discloses that CPLD 404, 604 is responsive to a set of input signals (address signal Add(b) and chip select signals cs0, cs1) to determine (“selectively electrically coupling”) an active rank of the four ranks and inactivating the other three ranks of memory devices from the computer system. *Id.* at 31–32 (citing Ex. 1008 ¶¶ 43, 44, 62); Ex. 1007 ¶ 66. According to Petitioner, the act in Amidi of activating one rank while deactivating other ranks constitutes “selectively coupling” and “selectively isolating.” *Id.* at 31–32.

Similarly, Petitioner cites to Ludwig’s disclosure of tri-state buffers in the VIC chip that isolate the loads of memory devices that are not enabled, to support Petitioner’s contention that Ludwig “selectively isolates one or

IPR2014-01011
Patent 7,881,150 B2

more loads of the DDR memory devices from the computer system.” *Id.* at 32.

According to Petitioner, a person of ordinary skill in the art would have been motivated to combine the teachings of Ludwig and Amidi, because (1) both references relate to a memory module including circuitry that is configurable to control one or more sets of memory devices based on signals provided by a connected computer and (2) both references are directed to solving the same problem of using additional memory devices in a memory module to include the overall memory density of the memory module without hindering the ability of the memory module to interface with a pre-existing memory controller. *Id.* at 39–40 (citing Ex. 1011, 2:10–22, 3:65–4:15; Ex. 1008 ¶¶ 38, 39, 43, 44, 62; Ex. 1009 ¶¶ 9,10,28); Ex. 1007 ¶¶ 116–119.

Petitioner supports its position with the Declaration of Dr. Jagannathan, who testifies that a person of ordinary skill in the art would have understood that the teachings of Amidi could have been advantageously applied to Ludwig. Ex. 1007 ¶ 117. In particular, Dr. Jagannathan opines that the teachings of Ludwig and Amidi could have been combined “to have more than one rank of memory devices such as the DDR SDRAM memories of Amidi, that can include more than one bank of memory arrays inside the DRAM device, and receive one of the chip select signals obtained by translating the input chip select and address signals.” *Id.* According to Dr. Jagannathan, the benefit of this combination is to apply the teachings of Ludwig to DDR SDRAM memories having faster access times. *Id.* Dr. Jagannathan concludes that one of ordinary skill would have been motivated to combine the architecture in Amidi with the VIC chip of

IPR2014-01011
Patent 7,881,150 B2

Ludwig to yield a circuit mounted in a memory module as disclosed by both Amidi and Ludwig. *Id.* ¶ 119.

Patent Owner contests Petitioner's position, arguing that the combination of Ludwig and Amidi fails to teach or suggest a "circuit configured to be mounted on a memory module[,] . . . the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal to the common data signal line," as recited in the challenged independent claims. PO Resp. 24. Patent Owner specifically argues that the Petition is deficient because the Petition's mapping of "selectively electrically coupling" to buffers 113 and 114 in Ludwig's Fig. 8 was incorrect. *Id.* at 25. According to Patent Owner, by equating the "selectively electrically coupling" to the buffers, Petitioner identified the "selectively electrically coupling" performed in Ludwig as being accomplished via circuitry internal to a memory chip (*i.e.*, via buffers 113 and 114) and not by the identified "circuit" (*i.e.*, VIC chip 90). *Id.* at 26. Patent Owner argues, however, claims 15 and 31 require the "circuit," not the "DDR memory device," to perform the "selectively electrically coupling." *Id.* Patent Owner further argues that the Petition fails to identify a "second data line" in Ludwig as required by the claims. *Id.*

We are unpersuaded by Patent Owner's position, because the arguments presented by Patent Owner generally attack the references individually, rather than in combination. Nonobviousness cannot be established by attacking the references individually when a challenge is predicated upon a combination of prior art disclosures. *See In re Merck & Co., Inc.*, 800 F.2d 1091, 1097 (Fed. Cir. 1986). In attacking the references

IPR2014-01011
Patent 7,881,150 B2

individually, Patent Owner fails to address Petitioner's actual challenges and establish an insufficiency in the combined teachings of the references. Despite Patent Owner's focus on Ludwig, we note that the Petition challenges the patentability of the claims based on the combination of Ludwig and Amidi. *See* Pet. 21–40. Based on Petitioner's arguments, we are satisfied that Amidi teaches or suggests : (i) a memory module with multiple ranks of DDR SDRAM memory devices 404 (Ex. 1008 ¶¶ 37, 42); (ii) CPLD 604 (*see id.* ¶ 28), register 608, and PLL 606 that can be implemented as a single component (*id.* ¶¶ 37, 50, Figs. 4A, 6A; *see* Ex. 1027, 38:1–39:15, 41:1–6); and (iii) electrically coupling CPLD 604, PLL 606, and register 608 to the ranks of DDR memory devices 404 (*id.* at Figs. 4A, 4B, 6A). We also credit the testimony of Dr. Jagannathan, who explains that “[w]hen Amidi's CPLD provides a chip select signal to a rank of memory devices, it selects the rank and thereby causes it to be coupled to the data bus.” Ex. 1028 ¶ 32 (citing Ex. 1008 ¶ 62). Thus, we are persuaded by Petitioner that Ludwig in combination with Amidi discloses or suggests a “circuit configured to be mounted on a memory module[,] . . . the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal to the common data signal line,” as recited in the challenged independent claims.

Petitioner further argues that Petitioner fails to identify a disclosure in Figure 8 of Ludwig of a “second data signal line” required by the challenged independent claims. PO Resp. 26. Patent Owner's position, however, appears to be premised on selected embodiments (specifically Figure 8) from Ludwig, rather than Ludwig's disclosure as a whole. Instead, we are persuaded by Petitioner's contention that certain embodiments from Ludwig

IPR2014-01011
Patent 7,881,150 B2

teach read and write lines DQ1–DQ4, which qualify as “first data signal lines” and “second data signal lines.” Pet. 24–25 (citing Ex. 1011, 9:31–40; Fig. 11).

Patent Owner also contends that under its proposed construction of “selectively electrically coupling,” neither Ludwig nor Amidi meets the claim limitation. PO Resp. 26–28. According to Patent Owner, the teachings of Ludwig are more similar to the generation of a signal, while Amidi is more similar to the transmitting of signals. *Id.* at 27. Patent Owner, then argues that “[t]he ’150 Patent discloses that ‘electrically couple’ is provided by structure (*e.g.*, “a plurality of electrical conduits which electrically couple” (Ex. 1001, 5:29–30), not by Ludwig’s signal generation and not by Amidi’s signal transmission. PO Resp. 27–28. Patent Owner explains that “[a]n electrical conduit in the ’150 Patent is a structural pathway for electricity, but the cited signals from Ludwig and Amidi are flows of electricity” (Ex. 2002 ¶ 79) and “[a] pathway for electricity (as for “selectively electrically coupling” in the ’150 Patent) is not a flow of electricity (as for the cited signals from Ludwig and Amidi) (*id.*).” PO Resp. 28.

Patent Owner then argues that the combination of Ludwig and Amidi fails to disclose “selectively electrically coupling” in the context of “signal line” as recited in claims 15 and 31. *Id.* at 28–29. According to Patent Owner, (i) such “signal line” structural elements are wholly missing from the teachings of Ludwig, because “memory layers and memory devices are not ‘signal line’ structural elements” and (ii) the data lines of Amidi are hard-wired in permanent coupling to the same data bus and are never subject to acts of “selectively electrically coupling.” *Id.* at 29–30. Patent Owner,

IPR2014-01011
Patent 7,881,150 B2

thus, concludes that “selectively electrically coupling” as in the ’150 Patent is distinct from the teachings of Ludwig and Amidi. *Id.* at 31.

As discussed previously, we do not agree with Patent Owner’s narrow construction of “selectively electrically coupling.” *See supra*, Section II.A.2. Rather, we construe the term as “making a selection between at least two components so as to transfer power or signal information from one selected component to at least the other selected component,” because such a construction is consistent with the ordinary and customary meaning of “selectively electrically coupling.” *Id.* Given our claim construction, we are unpersuaded by Patent Owner’s argument that the flow of electricity on hard-wired, permanent data signal lines does not constitute the electrical coupling of two selected components.

We also credit the testimony of Dr. Jagannathan, who states that Ludwig discloses VIC receiving chip-select signals and address signals as “selectively electrically coupling.” Ex. 1007 ¶ 115 (citing Ex. 1011, 7:16–30; Figs. 8, 10). According to Dr. Jagannathan, in response to one chip-select signal and two address signals, the VIC then generates chip-select signals that selectively enable one of four memory layers and selectively isolate the other memory devices in other ranks. We are persuaded that Ludwig’s teachings of “selectively enabling” and “selectively isolating” fall within the scope of the term “selectively electrically coupling” as we have construed the term. Furthermore, Amidi’s disclosure of directing signals down a specific signal line or data bus in order to determine an active rank within the memory devices falls within the scope of the term “selectively electrically coupling” as we have construed this term.

IPR2014-01011
Patent 7,881,150 B2

Patent Owner further contends that Ludwig and Amidi fail to disclose two separate “selectively electrically coupling” actions. PO Resp. 33. Patent Owner notes that claims 15 and 31 recite “the circuit . . . selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line.” *Id.* Patent Owner, thus, argues that the claims require that “the circuit” perform two separate acts of “selectively electrically coupling”: 1) a first act of “selectively electrically coupling” the first data signal line to the common signal line; and 2) a second act of “selectively electrically coupling” the second data signal line to the common signal line. *Id.* According to Patent Owner, Ludwig’s tri-state buffers 180 cannot “selectively electrically couple” a first data signal line to a common data signal line, and separately “selectively electrically couple” a second data signal line to the common data signal line because any selective electrical coupling that tristate buffers 180 may perform is performed collectively on both the first data signal line and the second data signal line. *Id.* at 38 (citing Ex. 2002 ¶ 82). As such, Patent Owner argues that the tri-state buffers 180 (and thus VIC chip 170) in Ludwig cannot perform two separate acts of “selectively electrically coupling” to couple a first data signal line and a second data signal line to a common data signal line, as required in challenged claims 15 and 31. *Id.* Patent Owner further argues that Amidi fails to make up for the deficiency of Ludwig, because Amidi simply discloses various signals and determining an active rank, and signals are not lines. *Id.* at 39.

We do not agree. To the contrary, we find that Ludwig teaches separate coupling actions. Specifically, Ludwig Figure 8 discloses that each

IPR2014-01011
Patent 7,881,150 B2

data signal line DQ1A-D has a respective data buffer to couple a corresponding memory chip from Layers A-D to common data bus DQ1, and each data buffer has a control logic that determines which of the four chips within the stack (Layers A-D) to couple. *See* Ex. 1011, 7:33–43, 7:56–65, Fig. 8. We further credit the testimony of Dr. Jagannathan, who states

Ludwig does disclose multiple acts of coupling. This is because Ludwig discloses selecting the active chip thereby causing selectively electrically coupling the data signal line of the active chip to the common data signal line. A [person of ordinary skill in the art] would understand that when a different active chip is selected, a different act of selectively electrically coupling is performed.

Ex. 1028 ¶ 55.

Finally, Patent Owner contends that Ludwig and Amidi are improperly combined by Petitioner. PO Resp. 39. Patent Owner explains that the Petitioner's combination is based on unsupported speculation and would result in a malfunctioning and inoperable system requiring heady redesign beyond the skill of a person of ordinary skill in the art. *Id.* at 42–43. Specifically, Patent Owner argues that Ludwig's older VIC interface chip signals are fundamentally different from the advanced DDR SDRAM technology disclosed in Amidi, and therefore, would not have been compatible or interchangeable. *Id.* at 40–43. According to Patent Owner, Amidi' DDR SDRAM technology is not backward compatible with the asynchronous DRAM technology of Ludwig, thus, the two technologies would be an inoperable combination. *Id.* at 39–40. Patent Owner specifically argues that in order to communicate properly with Amidi's DDR SDRAM devices, Ludwig's VIC chip would need to be modified to input

IPR2014-01011
Patent 7,881,150 B2

and output DDR command-based signals, requiring a redesign of the “fundamental operating principles of Ludwig’s VIC chip.” *Id.* at 42–43. According to Patent Owner, such a redesign would “undesirably gut the original design of Ludwig’s VIC chip.” *Id.* at 43.

Patent Owner relies on the Declaration of Dr. Sechen to support its position. Dr. Sechen testified that

A [person of ordinary skill in the art] would understand that such a redesign of Ludwig’s VIC chip and its components would change their fundamental operating principle from asynchronous to synchronous operation. Such a fundamental change would require a significant amount of complexity and design work to achieve a successful transformation in the fundamental aspect of timing. For example, where Ludwig’s VIC chip relies on its input and output signals to convey timing information by their own asynchronous waveforms, such timing information has to be conveyed through an entirely different mechanism in a synchronous environment where DDR commands are not designed to convey such timing information. A [person of ordinary skill in the art] might imagine that such a mechanism could be embodied into more logic, but it would not be readily apparent from Ludwig or Amidi how to implement such additional logic, thus requiring undue experimentation.

Ex. 2002 ¶ 93.

Despite Dr. Sechen’s explanation of asynchronous and synchronous memory (*see id.* ¶¶ 87–92), we do not agree that the teachings of Ludwig are incompatible with the synchronous memory of Amidi, because Ludwig expressly suggests applying its VIC chip to synchronized memory. *See Ex.* 1001, 10:10–13 (“A valuable added function might be synchronized memory signals to enhance the speed of the memory.”). We also credit the testimony of Dr. Jagannathan, who states

IPR2014-01011
Patent 7,881,150 B2

“[t]here was sufficient motivation for one of ordinary skill to apply the circuit including a CPLD, a PLL, and a register in Amidi to the interface chip in Ludwig. One of ordinary skill would recognize the benefits of doing so as taught in Ludwig. For example, Ludwig teaches ‘a valuable added function [of the VIC chip] might be synchronized memory signals to enhance the speed of the memory function.’”

Ex. 1007 ¶ 118 (citing Ex. 1011, 10:11–13). Therefore, we are persuaded that the teachings of Ludwig and Amidi are combinable.

Based on the evidence of record, we agree with Petitioner’s position that challenged independent claims 15, 22, and 31 would have been obvious over Ludwig and Amidi. Specifically, we are persuaded by Petitioner’s reasoning and the evidentiary record that Amidi teaches a circuit mounted on a memory module that is electrically coupled to a first DDR memory device and a second DDR memory device. We are further persuaded that Amidi teaches a circuit with a logic element, a register, and a PLL. We also are persuaded that the teachings of Amidi could have been implemented using the common data signal line and tri-state buffer system disclosed in Ludwig so that by (i) selectively electrically coupling a first data signal line to the common data signal line and (ii) selectively electrically coupling a second data signal line to the common data signal line, the circuit is responsive to a set of input signals. Additionally, we credit the testimony of Dr. Jagannathan that a person of ordinary skill in the art would have had reason to combine the teachings of Ludwig and Amidi, which both relate to memory modules, and are directed to solving the same problem of using additional memory devices in a memory module to increase the overall memory density of the module without hindering the ability of the module to interface with a pre-existing memory controller. *See id.* ¶ 116.

IPR2014-01011
Patent 7,881,150 B2

Additionally, we note that the testimony of Patent Owner’s Declarant, Dr. Sechen, is based on the claim constructions proffered by Patent Owner and not on the constructions set forth in the Decision to Institute. *Compare* Ex. 2002 ¶¶ 38–72, *with* Dec. to Inst. 7–12. We have considered as relevant, however, the portions of his analysis regarding the prior art and alleged non-obviousness of the claims and accorded the appropriate weight to that particular testimony.

Accordingly, we hold that Petitioner has shown by a preponderance of the evidence that claims 15, 22, and 31 of the ’150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Ludwig and Amidi.

b. Ludwig and Amidi Teach or Suggest All the Recited Limitations of Dependent Claims 16, 17, 24, 26, 32, and 33

Claims 16, 24, and 32 recite “wherein the two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.” Ex. 1001, 43:3–5, 44:1–3, 44:58–60. Petitioner contends the combined disclosures of Ludwig and Amidi, as summarized above, teach or suggest each limitation of dependent claims 16, 17, 24, 26, 32, and 33 of the ’150 patent. Pet. 38–40. Petitioner contends that Ludwig discloses that the VIC chip has functions for buffering, decoding (‘logic element’), refreshment for memory retention (‘register’), and synchronized memory signals (‘PLL’). *Id.* at 38 (citing Ex. 1004, 10:4–13). Petitioner further contends that Amidi discloses that the CPLD (“logic element”), register, and PLL are all mounted on a memory module. *Id.* (citing Ex. 1008 ¶¶ 37, 39, 40, Figs. 4A, 4B, 6). According to Petitioner, a person of ordinary skill in

IPR2014-01011
Patent 7,881,150 B2

the art would have recognized that at least the register and PLL could be portions of a single component. *Id.* (citing Ex. 1007 ¶¶ 58, 108).

Dependent claim 17 recites that the circuit includes “one or more switches selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the one or more switches operatively coupled to the logic element to receive control signals from the logic element.” Dependent claim 33 recites a similar limitation. Ex. 1001, 43:6–12. Petitioner contends Ludwig meet the limitations of claims 17 and 33, because a person of ordinary skill in the art would have recognized the tri-state buffers in Ludwig are “operatively coupled” to the logic element to receive control signals from the logic element. Pet. 39 (citing Ex. 1011, 7:33–45, 9:37–39, Figs. 7, 8, 11; Ex. 1007 ¶¶ 105, 112).

Dependent claim 26 further recites that the claimed circuit is “configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.” Ex. 1001, 44:8–11. Petitioner contends Ludwig discloses that decoder 64 (“circuit”) enables one of four memory chips 60 and disenables (“selectively isolate[s]”) the other three memory chips (“a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system”). Pet. 39 (citing Ex. 1011, 6:15–20). Petitioner further contends that Amidi teaches that the PLL and register isolate the loads of the memory devices that are not enabled by the chip select signal generated by the CPLD. *Id.* at 39–40 (citing Ex. 1008 ¶¶ 38–39, 43, 62, Figs. 4A and 5). According to Petitioner, a person of ordinary skill in the art would have recognized that selectively isolating the loads of the memory devices would

IPR2014-01011
Patent 7,881,150 B2

require isolating the data signal lines of the memory devices from the computer system. *Id.* at 40 (citing Ex. 1007 ¶¶ 72, 105).

Patent Owner does not provide separate contentions regarding the limitations recited in the dependent claims, but relies on the arguments made in support of patentability of independent claims 15, 22, and 31. *See generally* PO Resp.

After consideration of the language recited in claims 16, 17, 24, 26, 32, and 33 of the '150 patent, the Petition, the Patent Owner Response, and Petitioner's Reply, as well as the relevant evidence discussed in those papers, we find that one of ordinary skill in the art would have considered these dependent claims obvious over the combination of Ludwig and Amidi. Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that claims 16, 17, 24, 26, 32, and 33 of the '150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Ludwig and Amidi.

F. Asserted Anticipation of Claims 22, 24, and 26 in view of Amidi

Petitioner alleges claims 22, 24, and 26 of the '150 patent are unpatentable under 35 U.S.C. § 102(e) in view of Amidi. Pet. 40. Patent Owner disputes Petitioner's position, arguing that Amidi fails to disclose all of the claim limitations. PO Resp. 24–26, 32–39.

We have reviewed the Petition, the Patent Owner Response, and Petitioner's Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 22, 24, and 26 of the '150 patent are anticipated by Amidi.

IPR2014-01011
Patent 7,881,150 B2

1. Overview of Amidi

See Section II.D.2., discussed above.

2. Analysis

Petitioner contends Amidi, as summarized above, discloses each limitation of claims 22, 24, and 26 of the '150 patent. Pet. 40 (citing *id.* at 21–40). Petitioner first argues that Amidi discloses a circuit mounted on a memory module, where the circuit includes a logic element, a register, and a phase-lock loop device. *Id.* at 22 (citing Ex. 1008 ¶¶ 2, 37, Figs. 4A, 6A); Ex. 1007 ¶¶ 58, 64. Petitioner also argues that the system described in Amidi includes a memory module having one or more ranks of DDR memory devices, which are electrically coupled to the components of a circuit. *Id.* at 22–23 (citing Ex. 1008, Figs. 4A, 4B, 6A.) Petitioner supports its position with the declaration of Dr. Jagannathan, who testifies that Amidi teaches a circuit that includes a CPLD, a register, and a PLL circuit. Ex. 1007 ¶ 58. Dr. Jagannathan further testifies that “[o]ne of ordinary skill would understand that these electrical components or devices are electrically coupled together to form a ‘circuit.’” *Id.* (citing Ex. 1008 ¶¶ 34–39, 41).

Petitioner then argues that Amidi also discloses memory devices having data signal lines and data strobe lines. *Id.* at 27 (citing Ex. 1008 ¶¶ 29, 32, Figs. 2, 3). According to Petitioner, Amidi discloses that each stack of DDR memory devices has a data signal line and a data strobe line DQS. (*id.* at 26 (citing Ex. 1008 ¶ 32, Fig. 3)), and that at least two DDR memory devices are connected to the same (common) memory bus (“common data signal line”) (*id.* (citing Ex. 1008 ¶¶ 34–35, Fig. 3)). According to Petitioner, each DDR memory device has its own data pins (data bus) that are connected to a common data signal line, and the circuit of Amidi is

IPR2014-01011
Patent 7,881,150 B2

“electrically coupled” to the common data bus. *Id.* at 26; *see* Ex. 1007 ¶¶ 63, 72.

Petitioner provides arguments for the sending and receiving of input signals by the system in Amidi, and explains how Amidi discloses a circuit that is responsive to such input signals. *Id.* at 27–38. Petitioner also argues that Amidi teaches “selectively coupling” because Amidi discloses that CPLD 404, 604 is responsive to a set of input signals (address signal Add(b) and chip select signals cs0, cs1) to determine (“selectively electrically coupling”) an active rank of the four ranks and inactivating the other three ranks of memory devices from the computer system. *Id.* at 31–32 (citing Ex. 1008 ¶¶ 43, 44, 62); Ex. 1007 ¶ 66. According to Petitioner, the act in Amidi of activating one rank while deactivating other ranks constitutes “selectively coupling” and “selectively isolating.” *Id.*

Petitioner then contends Amidi discloses that “memory module 400 includes 92 contact pins 402 on the front side for connecting with a memory socket,” which is a system memory domain. *Id.* at 37 (citing Ex. 1008 ¶ 37). According to Petitioner, Amidi discloses that module connector 602 (contact pins 402 and 416) sends cs0 and cs1 signals, which can be considered “first number of chip-select signals,” to CPLD 604. *Id.*

Patent Owner contests Petitioner’s position, arguing that Amidi fails to disclose a “circuit . . . selectively isolating one or more loads of the DDR memory devices from the computer system,” as recited in the challenged claims. PO Resp. 46–50. According to Patent Owner, “choosing a rank of memory devices and inactivating other ranks is not selectively isolating a load of DDR memory circuits, because neither involves electrical separation from the computer system.” *Id.* at 47. Patent Owner contends that Amidi

IPR2014-01011
Patent 7,881,150 B2

only teaches hard-wired lines in permanent connection to the same 72-line data bus. *Id.* at 47–48 (citing Ex. 1008 ¶¶ 34, 35, Fig. 3). Patent Owner supports its position with the declaration of Dr. Sechen, who testifies that

the data bus lines from the memory controller at all times are electrically coupled to the memory devices. That’s because there are no switches on the chip select signal lines nor the data bus lines to make possible any type of electrical decoupling.

Ex. 2002 ¶ 78. Patent Owner, thus, concludes that due to the permanent electrical connections of Amidi’s data bus, there would never be electrical separation and, thus, the memory devices could never be subject to acts of “selectively isolating.” PO Resp. 48.

Patent Owner further contends the limitation “selectively isolating one or more loads” is distinct from the term “selectively electrically coupling” and is not disclosed by Amidi. *Id.* at 46–47. Patent Owner argues that the rank activation and inactivation in Amidi is “some kind of operational isolation” but is not a load isolation within the scope of the claims. *Id.* at 49.

We do not agree with Patent Owner’s position that hard-wired data signal lines cannot be electrically isolated in a selective fashion. As discussed above, we construe “selectively electrically isolating” as “making a selection between at least two components and not transferring power or signal information from one selected component to the other selected component.” *See supra*, Section II.A.2. Amidi’s disclosure of directing electrical signals down a specific signal line or data bus in order to electrically activate a rank within the memory devices and electrically inactivate other ranks falls within the scope of the term “selectively electrically isolating” as we have construed the term. We are further unpersuaded by Patent Owner’s argument that “selectively isolate” or

IPR2014-01011
Patent 7,881,150 B2

“selectively electrically isolating” is limited narrowly to a particular kind of load isolation. The ’150 patent defines the term “load” broadly and includes “*without limitation*, electrical load, such as capacitive load, inductive load, or impedance load.” Ex. 1001, 5:3–5 (emphasis added). Thus, Patent Owner’s arguments are not commensurate with the ’150 patent definition of “load” or with the scope of the challenged claims.

Based on the evidence of record, we agree with Petitioner’s position that challenged claims 22, 24, and 26 are anticipated by Amidi. First, we are persuaded by Petitioner’s reasoning and the evidentiary record that Amidi discloses the recited limitations of claim 22. *See* Ex. 1001, 43:35–61. Specifically Amidi discloses: (i) a circuit with a logic element, a register, and a PLL (*see* Ex. 1008 ¶¶ 34–39, 41, Fig. 4A); (ii) a circuit mounted on a memory module that is electrically coupled to a plurality of DDR memory devices arranged in ranks on a memory module (*see id.* ¶¶ 4, 12, 34–39, Fig. 4A); (iii) the memory module electrically coupled to a memory controller (*see id.* ¶¶ 29–32, 41); (iv) the memory module received a set of input signals (*see id.* ¶¶ 29–32, 50, 58, Fig. 3); (v) the circuit responsive to the input signals by selectively isolating one or more loads of the DDR memory devices from the computer system (*see id.* ¶¶ 38, 39, 43, 44, 62, Fig. 6A); and (vi) the circuit translates between the system memory domain of the computer system and a physical memory domain of the plurality of the DDR memory devices (*see id.* ¶¶ 34, 35, 37, 41–43, 49, 50, 52, 57, 60, 62, Figs. 5, 7).

Second, we are persuaded by Petitioner’s reasoning and the evidentiary record that Amidi discloses the additional recited limitations of claim 26. Dependent claim 26 further recites that the claimed circuit is

IPR2014-01011
Patent 7,881,150 B2

“configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.” Ex. 1001, 44:8–11. Amidi teaches that the PLL and register isolate the loads of the memory devices that are not enabled by the chip select signal generated by the CPLD. *See* Ex. 1008 ¶¶ 38–39, 43, 62, Figs. 4A, 5.

According to Petitioner, a person of ordinary skill in the art would recognize that selectively isolating the loads of the memory devices would necessarily require isolating the data signal lines of the memory devices from the computer system. Pet. 39–40 (citing Ex. 1007 ¶¶ 72, 105). We agree.

We are not persuaded, however, by Petitioner’s reasoning and the evidentiary record that Amidi discloses the additional limitations of claim 24 as specifically recited and arranged in claim 24. Dependent claim 24 recites “wherein the two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.” Ex. 1001, 44:1–3. Amidi discloses that the CPLD (“logic element”), register, and PLL are all mounted on a memory module. Ex. 1008 ¶¶ 37, 39, 40, Figs. 4A, 4B, 6). According to Petitioner, a person of ordinary skill in the art would recognize that at least the register and PLL could be portions of a single component. Pet. 38 (citing Ex. 1007 ¶¶ 58, 108). Amidi, however, does not disclose specifically, nor is it inherent in Amidi, that two or more of the logic element, the register, and the phase-lock loop device mounted on a single memory module are portions of a single component. Therefore, we are not persuaded Petitioner has carried its burden to demonstrate, by a preponderance of the evidence, that Amidi anticipates claim 24 of the ’150 patent.

IPR2014-01011
Patent 7,881,150 B2

Accordingly, we hold that Petitioner has shown by a preponderance of the evidence that claims 22 and 26 of the '150 patent are unpatentable under 35 U.S.C. § 102(e) as anticipated by Amidi.

III. CONCLUSION

We conclude Petitioner has shown by a preponderance of the evidence that (1) claims 15–17, 22, 24, 26, and 31–33 of the '150 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combinations of Ludwig and Amidi and (2) claims 22 and 26 of the '150 patent are unpatentable under 35 U.S.C. § 102(e) as anticipated by Amidi.

IV. ORDER

For the reasons given, it is

ORDERED that, by a preponderance of the evidence, claims 15–17, 22, 24, 26, and 31–33 of the '150 patent are unpatentable; and

FURTHER ORDERED that because this is a Final Written Decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2014-01011
Patent 7,881,150 B2

FOR PETITIONER:

Fabio Marino
fmarino@mwe.com

Bernie Knight
bknight@mwe.com

Barrington Dyer
bdyer@mwe.com

MCDERMOTT WILL & EMERY

FOR PATENT OWNER:

Mehran Arjomand
marjomand@mofo.com
patentdocket@mofo.com

David Kim
dkim@mofo.com
patentdocket@mofo.com

Erol Basol
ebasol@mofo.com
patentdocket@mofo.com

Jean Nguyen
jnguyen@mofo.com
patentdocket@mofo.com

Jonathan Statman
jstatman@mofo.com
patentdocket@mofo.com

MORRISON & FOERSTER LLP

Trials@uspto.gov
571.272.7822

Paper No. 33
Filed: December 14, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DIABLO TECHNOLOGIES, INC.,
Petitioner,

v.

NETLIST, INC.,
Patent Owner.

Case IPR2014-00883
Patent 8,081,536 B1

Before LINDA M. GAUDETTE, BRYAN F. MOORE, and
GEORGIANNA W. BRADEN, *Administrative Patent Judges*.

BRADEN, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318 and 37 C.F.R. § 42.73

IPR2014-00883
Patent 8,081,536 B1

I. INTRODUCTION

We have jurisdiction to hear this *inter partes* review under 35 U.S.C. § 6(c). This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 1, 16, 17, 24, 30, and 31 of U.S. Patent No. 8,081,536 B1 (Ex. 1001, “the ’536 patent”) are unpatentable.

A. Procedural History

Diablo Technologies, Inc. (“Petitioner”) filed a Corrected Petition (Paper 4, “Pet.”) to institute an *inter partes* review of claims 1, 16, 17, 24, 30, and 31 of the ’536 patent. Netlist, Inc. (“Patent Owner”) filed a Preliminary Response (Paper 8, “Prelim. Resp.”). Pursuant to 35 U.S.C. § 314(a), we instituted an *inter partes* review of all challenged claims on the following grounds alleged in the Petition.

References	Basis	Claims Challenged
Klein ¹ and Amidi ²	§ 103	1, 16, 17, 24, 30, and 31
Klein, Amidi, and Dell ³	§ 103	16, 17, 30, and 31

Paper 11 (“Dec. to Inst.”), 26.

After institution of trial, Patent Owner filed a Patent Owner Response (Paper 25, “PO Resp.”), to which Petitioner filed a Reply (Paper 27, “Reply”). An oral argument was held on July 28, 2015, consolidated with

¹ US Patent Publication No. 2001/0008006 A1, pub. July 12, 2001 (“Klein,” Ex. 1006).

² US Patent Publication No. 2006/0117152 A1, pub. June 1, 2006 (filed Jan. 5, 2004) (“Amidi,” Ex. 1008).

³ U.S. Patent No. 6,446,184 B2, iss. Sept. 3, 2002 (“Dell,” Ex. 1009).

IPR2014-00883
Patent 8,081,536 B1

the oral hearings for IPR2014-00882 and IPR2014-01011. *See* Paper 30. A transcript (“Tr.”) of the oral argument is included in the record. Paper 31.

B. Related Proceedings

The parties inform us that the ’536 patent is the subject of the following federal district court case: *Netlist, Inc. v. Smart Modular Technologies*, Case. No. 4:13-cv-05889-YGR (N.D. Cal.). Papers 6, 10.

The ’536 patent claims priority to U.S. Patent No. 7,289,386 (“the ’386 patent). Pet. 7. Petitioner informs us that the related ’386 patent is the subject of *inter partes* reexamination 95/000,577. *Id.* The ’386 patent is also the subject of district court case *Google, Inc. v. Netlist, Inc.*, Case No. C08-4144 SBA (N.D. Cal.). *Id.* at 11.

In addition, Petitioner filed two other petitions requesting *inter partes* review of related U.S. Patent No. 7,881,150 B2. Paper 6, 2. These cases are: IPR 2014-00882 and IPR2014-01011. *Id.* We consolidated the oral hearings for IPR2014-00882, IPR2014-00883, and IPR 2014-01011. *See* Paper 30.

C. The ’536 Patent

The ’536 patent is directed to a memory module of a computer system with improved performance and memory capacity. Ex. 1001, 1:35–38. Memory module 10 includes a plurality of memory devices 30 (arranged in ranks 32) and circuit 40. *Id.* at 5:25–34; Fig. 1. Circuit 40 is electrically coupled to the memory devices 30 and memory controller 20 of a computer system. *Id.* The memory module improves performance and memory capacity by isolating electrical loads of the memory devices from the computer system. *Id.* at 5:34–35.

IPR2014-00883
Patent 8,081,536 B1

Circuit 40 receives input signals from memory controller 20. *Id.*

Figure 1, reproduced below, illustrates input signals (corresponding to a number of memory devices) from memory controller 20, such as chip select signals (“cs#”), that are directed to memory module 10, which can act as a virtual memory module. *Id.* at 16:57–65; Figs. 1, 9A, 9B.

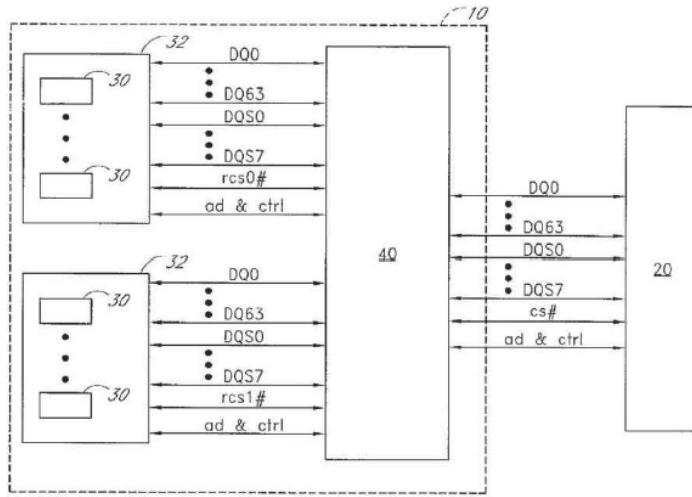


Figure 1 is a schematic of a memory module with circuit 40 and memory devices 30 and connectable to memory controller 20.

As shown in Figure 1, based on the received input signals from memory controller 20, circuit 40 generates output signals corresponding to memory devices 30 on the memory module. *Id.* at Fig. 1. The output signals include a different number of chip select signals (e.g., “rcs0#” and “rcs1#”) corresponding to memory devices 30 shown in ranks 32. *Id.* at 17:7–12; Fig. 1.

Circuit 40 includes a logic element, such as a programmable logic device (PLD), an application-specific integrated circuit (ASIC), a field programmable gate array (FPGA), or a complex programmable-logic device (CPLD). Ex. 1001, 6:40–45. As shown in Figure 9A, reproduced below,

IPR2014-00883
Patent 8,081,536 B1

circuit 40 also includes register 230 and phase-lock loop device (PLL) 220. *Id.* at 15:52–58; Fig. 9A.

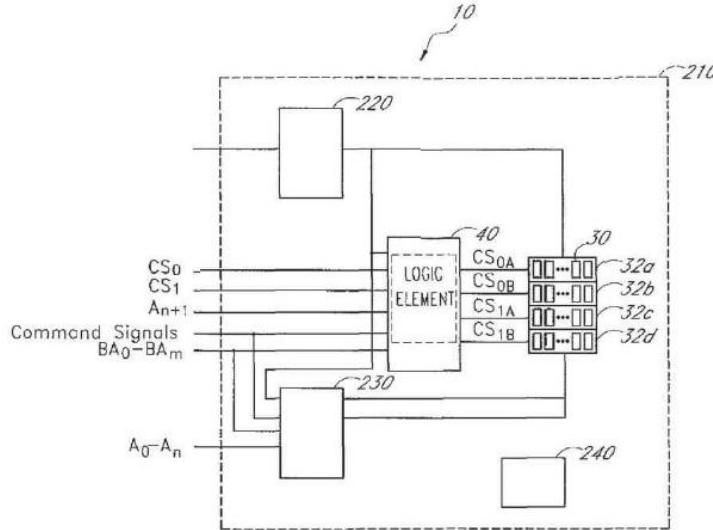


Figure 9A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Figure 9A illustrates circuit 40 receiving a set of input command signals, address signals (A_{n+1}), including bank address signals (BA₀-BA_m), row address signals (A₀-A_n), column address signals, gated column address strobe signals, and chip-select signals (CS₀, CS₁), from memory controller 20 of the computer system. *Id.* at 16:36–41, 17:19–34. “In response to the set of input address and command signals, circuit 40 generates a set of output address and command signals.” *Id.* at 16:42–44.

With the output address and command signals, circuit 40 isolates the electrical loads of some of memory devices 30 from the computer system. *Id.* at 7:17–31. According to the '536 patent, load isolation may result in specific benefits including reduced load related to data signal lines. *Id.* at 14:57–62. In order to isolate the loads, the logic element of circuit 40 translates between a system memory domain of the computer system and a

IPR2014-00883
Patent 8,081,536 B1

physical memory domain of memory module 10. *Id.* at 7:17–31. As shown in Figure 3, reproduced below, the circuit isolates the load of a memory device by isolating one or both of DQ data signal lines 102a, 102b of two memory devices 30a and 30b from common DQ data signal line 112 that is coupled to the computer system. Ex. 1001, 7:32–38, Fig. 3A.

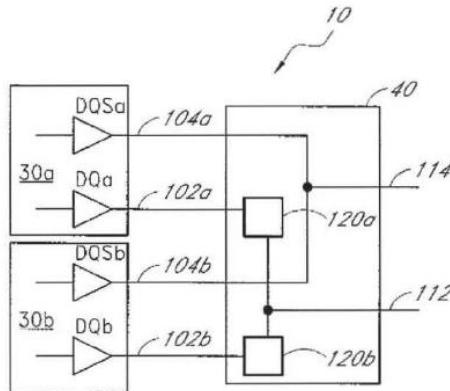


Figure 3A is a schematic of circuit 40 receiving a set of input command signals from memory controller 20 of the computer system.

Circuit 40 can electrically couple one or both of the DQ data signal lines 102a, 102b of the two memory devices 30a and 30b to the common data signal line 112, at the same time. *Id.* at 7:58–62. Circuit 40 also allows a DQ data signal to be transmitted from memory controller 20 of the computer system to one or both of DQ data signal lines 102a, 102b. *Id.* at 7:38–41. The logic element of circuit 40 uses switches 120a, 120b in order to isolate or couple one or both of DQ data signal lines 102a, 102b of memory devices 30a and 30b from common data signal line 112. *Id.* at 7:38–46.

D. Illustrative Claim

As noted above, *inter partes* review was instituted for claims 1, 16, 17, 24, 30, and 31 of the '536 patent, of which claims 1 and 24 are

IPR2014-00883
Patent 8,081,536 B1

independent claims. Claim 1 is illustrative of the challenged claims and is reproduced below:

1. A circuit configured to be mounted on a memory module configured to be operationally coupled to a computer system, the memory module having a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of DDR chip-select signals, the circuit including at least one configuration in which the circuit is configured to:
 - receive a set of signals comprising address signals and a second number of DDR chip-select signals smaller than the first number of DDR chip-select signals;
 - generate phase-locked clock signals and transmit the phase-locked clock signals to the DDR memory circuits of the first number of ranks;
 - selectively isolate a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals; and
 - generate the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of DDR chip-select signals.

Ex. 1001, 41:20–43.

II. DISCUSSION

A. *Claim Construction*

In an *inter partes* review, claim terms in an unexpired patent are interpreted according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b); *see also In re Cuozzo Speed Techs., LLC*, 793 F.3d 1268, 1278–79 (“Congress implicitly approved the broadest reasonable interpretation standard in enacting the AIA,” and “the standard was properly adopted by PTO

IPR2014-00883
Patent 8,081,536 B1

regulation.”). Under that standard, and absent any special definitions, we give claim terms their ordinary and customary meaning as would be understood by one of ordinary skill in the art at the time of the invention. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). Yet a “claim term will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.” *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002).

1. “Selectively Isolate” and “Selectively Isolating,”

In the Decision to Institute, we construed the terms “Selectively Isolate” and “Selectively Isolating,” which are recited in all the challenged independent claims. *See* Dec. to Inst. 8–10. In their papers, neither party challenged our constructions of these claim terms.⁴ PO Resp. 4, 25–27; Reply 3–8. Thus, for the terms “Selectively Isolate” and “Selectively Isolating,” we see no reason to alter the constructions of these claim terms as set forth in the Decision to Institute as shown below, and we incorporate our previous analysis for purposes of this decision.

⁴ Patent Owner’s Declarant, Dr. Sechen, testifies based on the claim constructions in Patent Owner’s Preliminary Response and not on the constructions set forth in the Decision to Institute. *Compare* Ex. 2002 ¶¶ 38–70 (Declaration of Dr. Carl Sechen), *with* Dec. to Inst. 8–10. In several instances, Dr. Sechen acknowledges that Patent Owner does not propose constructions in its Patent Owner Response. Ex. 2002 ¶¶ 38, 59, 64. Dr. Sechen does not contend we should modify our claim constructions, nor does he provide sufficient rationale as to why we should modify our claim constructions. *Id.*

IPR2014-00883
Patent 8,081,536 B1

Claim Term	Construction
“selectively isolate” / “selectively isolating”	“electrical separation from one selected component from another selected component”

See Dec. to Inst. 8–10.

Patent Owner, however, argued at the oral hearing for a different construction for “Circuit Configured to be Mounted on a Memory Module.” Tr. 68:1–18. Therefore, we address Patent Owner’s contentions and construe “Circuit Configured to be Mounted on a Memory Module” as discussed below.

2. *“Circuit Configured to be Mounted on a Memory Module”*

In the Decision to Institute, we construed the term “a circuit configured to be mounted on a memory module,” to encompass “circuitry configured to be mounted on at least a portion of a memory module.” Dec. to Inst. 11–12. Such a construction is consistent with the ordinary and customary meaning of “a circuit configured to be mounted on a memory module.” *Id.* at 12.

Petitioner agrees with the construction set forth in the Decision to Institute. Reply 4; Tr. 5:25–6:10. Patent Owner, however, contends that “a circuit configured to be mounted on a memory module” should be construed as “an entire circuit configured to be mounted on a single memory module.” PO Resp. 11–13; Tr. 68:1–18.

Patent Owner notes that our claim construction, as set forth in the Decision to Institute, is ambiguous in that it can be read two ways:

One might read the Board’s construction as meaning that the circuit is mounted on and occupies at least a portion of the memory module (Ex. 2002, ¶ 62), which may be consistent with

IPR2014-00883
Patent 8,081,536 B1

Netlist's construction. On the other hand, one might read the Board's construction as encompassing portions of the circuit to be mounted off-module, which would be unreasonably broad to a [person of ordinary skill in the art].

Id. at 11 (citing Ex. 2002 ¶ 62). According to Patent Owner, “memory module” is a term of art that would have had a well-understood meaning to a person of ordinary skill in the art at the time of the invention and a person of ordinary skill in the art would not have understood “a circuit configured to be mounted on a memory module” to include circuit parts off of that memory module or on a different memory module. *Id.* at 12 (citing Ex. 2002 ¶ 60). Patent Owner contends that such a construction could include situations that defeat the purpose of a memory module to make removing and installing memory upgrades easy and error-free. *Id.* (citing Ex. 2002 ¶ 61).

We decline to adopt Patent Owner’s claim construction as it is inconsistent with the definition of “circuit” as found in the specification of the ’150 patent. The ’150 patent defines “circuit” as “a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.” Ex. 1001, 5:9–13. The ’150 patent does not limit a “circuit” to only a configuration of electrical components or devices that are mounted on a single memory module. Therefore, applying the broadest reasonable interpretation consistent with the specification of the ’150 patent, we construe the claim element “a circuit configured to be mounted on a memory module,” as we did in the Decision to Institute, but we further

IPR2014-00883
Patent 8,081,536 B1

clarify the construction to encompass “at least a portion of circuitry configured to be mounted on at least a portion of a memory module.”

3. Other Claim Terms

We determine that no express constructions of any other claims terms are required for our analysis, and we apply the ordinary and customary meaning of each claim term.

B. Principles of Law

To prevail in its challenges to the patentability of the claims, a petitioner must establish facts supporting its challenges by a preponderance of the evidence. 35 U.S.C. § 316(e); 37 C.F.R. § 42.1(d). A claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.

KSR Int'l Co. v. Teleflex Inc., 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

We analyze the instituted grounds of unpatentability in accordance with the above-stated principles.

C. Level of Ordinary Skill in the Art

In determining whether an invention would have been obvious at the time it was made, we determine the level of ordinary skill in the pertinent art at the time of the invention. *Graham v. John Deere*, 383 U.S. at 17. “The

IPR2014-00883
Patent 8,081,536 B1

importance of resolving the level of ordinary skill in the art lies in the necessity of maintaining objectivity in the obviousness inquiry.” *Ryko Mfg. Co. v. Nu-Star, Inc.*, 950 F.2d 714, 718 (Fed. Cir. 1991).

Petitioner’s Declarant, Srinivasan Jagannathan, Ph.D. (“Dr. Jagannathan”), testifies that a person of ordinary skill in the art at the time of the ’536 patent:

would understand basic memory and data communication concepts, with a bachelor’s degree in any of electrical engineering, computer engineering, computer science, or related field. Course work for one of ordinary skill in the art would have included a course on computer organization, principles of digital design, or computer architecture. One of ordinary skill in the art would also have around one year of experience related to computer memory systems. For example, such experience may include experience in DRAM memory technology and related industry standards such as JEDEC standards for DRAM memories and memory modules.

Ex. 1011 ¶ 50.

Patent Owner’s Declarant, Carl Sechen, Ph.D. (“Dr. Sechen”), testifies that one of ordinary skill in the art at the time of the ’536 patent would have had an undergraduate degree in electrical engineering or computer engineering, at least two years of professional experience in the design of memory systems, familiarity with the latest JEDEC standard specifications for memory devices and modules, and familiarity with the latest DRAM memory devices widely available in the market. Ex. 2002 ¶ 14. Dr. Sechen further testifies that a person of ordinary skill in the art would have design proficiency in memory modules comprising DDR memory technology, such as memory modules with JEDEC standard DDR SDRAM devices. *Id.* at ¶ 15. Patent Owner concurs with Dr. Sechen’s

IPR2014-00883
Patent 8,081,536 B1

opinion regarding the level of skill in the art at the time of the '536 patent.
PO Resp. 13–15 (citing Ex. 2002 ¶¶ 14–16, 18–20, 28).

Based on our review of the '536 patent and the types of problems and solutions described in the '536 patent and cited prior art, we conclude a person of ordinary skill in the art at the time of the '536 patent would have a Bachelor's degree in electrical engineering, computer engineering, computer science, or related field and at least one year of work experience, including familiarity with computer memory systems and related industry standards such as JEDEC standards for DRAM memories and memory modules. We further note that the applied prior art reflects the appropriate level of skill at the time of the claimed invention. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001).

D. Expert Testimony

Patent Owner argues that Petitioner's declarant, Dr. Jagannathan, does not qualify as a person of ordinary skill in the art at the time of the invention because of his alleged lack of experience designing memory modules. PO. Resp. 13–22. According to Patent Owner, Dr. Jagannathan's experience and background is directed to software and is not relevant to the case. *Id.* at 15–20. Patent Owner argues that, unlike Dr. Jagannathan, its Declarant Dr. Sechen has significant practical experience designing memory modules. *Id.* at 20 (citing Ex. 2002, Exhibit A).

As to his hardware and memory design experience, Dr. Jagannathan was awarded a Doctorate degree in Computer Science, and has “two decades of experience in the design, development, and analysis of a wide range of hardware, software, network and database systems.” Ex. 1011 ¶ 2. He has designed and implemented hardware virtual memory caches, and researched

IPR2014-00883
Patent 8,081,536 B1

theoretical performance measures of various cache coherency protocols. *Id.* Dr. Jagannathan also stated during his deposition that he has experience in the design of memory systems. Ex. 2003, 124:12–126:1. Patent Owner contends, however, that Dr. Jagannathan fails to qualify as a person of ordinary skill in the art because he lacks sufficient professional experience physically designing (i.e., “actually putting down a design and saying this is what it would be”) a memory module. PO Resp. 18 (citing Ex. 2003, 125:14–17). We disagree. To testify as an expert under Fed. R. Evid. 702, a person need not be a person of ordinary skill in the art, but rather must be “qualified in the pertinent art.” *Sundance, Inc. v. DeMonte Fabricating Ltd.*, 550 F.3d 1356, 1363–64 (Fed. Cir. 2008); *see SEB S.A. v. Montgomery Ward & Co.*, 594 F.3d 1360, 1372–73 (Fed. Cir. 2010) (upholding a district court’s ruling to allow an expert to provide testimony at trial because the expert “had sufficient relevant technical expertise” and the expert’s “knowledge, skill, experience, training [and] education . . . [wa]s likely to assist the trier of fact to understand the evidence”); *Mytee Prods., Inc. v. Harris Research, Inc.*, 439 Fed. App’x 882, 886–87 (Fed. Cir. 2011) (non-precedential) (upholding admission of the testimony of an expert who “had experience relevant to the field of the invention,” despite admission that he was not a person of ordinary skill in the art). We find that, although Dr. Jagannathan is less experienced than Dr. Sechen in the area of memory module design, he is qualified sufficiently to testify as an expert witness about memory systems and memory modules.

IPR2014-00883
Patent 8,081,536 B1

E. Alleged Obviousness of Claims 1, 16, 17, 24, 30, and 31 in view of Klein and Amidi

Petitioner alleges claims 1, 16, 17, 24, 30, and 31 of the '536 patent are unpatentable under 35 U.S.C. § 103 in view of Klein and Amidi. Pet. 18–27. Patent Owner disputes Petitioner's position, arguing that a person of ordinary skill in the art would not have had reason to combine the references in the manner proposed by Petitioner (PO Resp. 28) and further that the combination of the references fails to teach or suggest all of the claim limitations (*id.* at 4–5, 22–28).

We have reviewed the Petition, the Patent Owner's Response, and Petitioner's Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 1, 16, 17, 24, 30, and 31 of the '536 patent are unpatentable as obvious over the combination of Klein and Amidi.

1. Overview of Klein

Klein discloses a method for bus capacitance reduction. Ex. 1006, Abstract. According to Klein, data bus capacitance is reduced by decoupling unaccessed memory circuits from a data bus during data transfers to or from other memory circuits. *Id.* One embodiment in Klein provides memory controller 22 connects to circuitry 26 for interfacing with one or more memory circuits 28, as shown in Figure 3, reproduced below. *Id.* ¶ 28, Fig. 3.

IPR2014-00883
Patent 8,081,536 B1

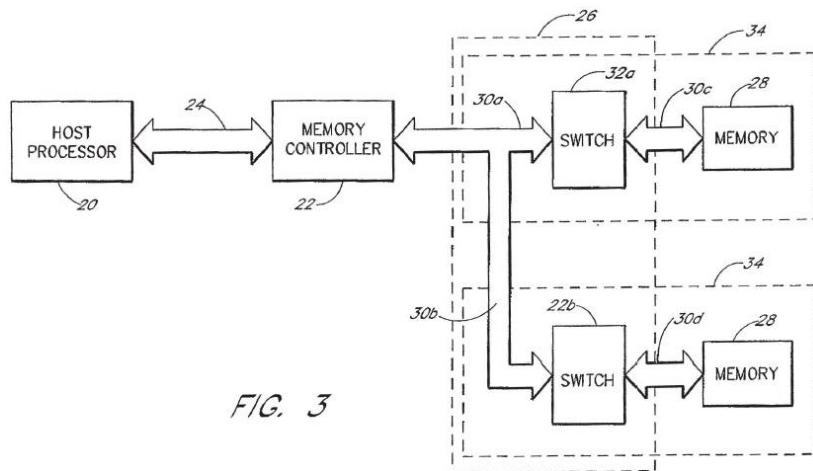


Figure 3 is a schematic of a bus switch that couples or decouples memory elements 28 and memory controller 22.

Figure 3 illustrates that the data bus between memory controller 22 and memory elements 28 may comprise several branches 30a, 30b, one for each separate memory element 28. *Id.* Each branch may include switch 32a, 32b, that may be used to selectively isolate portions (30c, 30d) of the data bus running from memory controller 22 to circuitry memory element 28. *Id.* Klein states that memory circuit 28 may be a conventional DRAM integrated circuit. *Id.* ¶ 29. According to Klein, the embodiment shown in Figure 3 may reduce the parasitic capacitance that the memory controller needs to charge and discharge during data transfers because a portion of the data bus and the stray capacitance of unaccessed memory circuits are removed. *Id.* ¶ 28.

Another embodiment in Klein is illustrated in Figure 6, reproduced below.

IPR2014-00883
Patent 8,081,536 B1

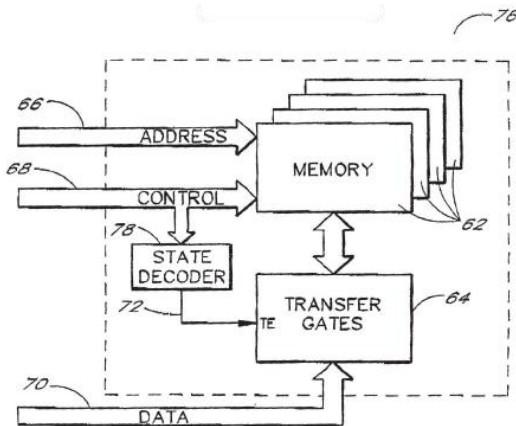


Figure 6 is a schematic of a memory module with memory elements that connect to an integrated circuit with transfer gates and state decoder.

As shown in Figure 6, a circuit is provided on memory module 76 that includes transfer gates 64 and state decoder 78. Ex. 1006 ¶¶ 35, 39. Klein discloses that state decoder 78 includes inverter 80 (*id. ¶ 36*), and that “the state decoder 78 could comprise a state machine 84 made with a programmable gate array for example” (*Id. ¶ 37*). Also, Klein teaches that the state decoder may be implemented as a state machine. *Id. ¶ 37*, Fig. 8.

Klein further discloses control logic circuitry, data buffer registers, and a bus switch that is incorporated into memory modules. *Id. ¶¶ 29, 39, 40*; Figs. 3, 10. According to Klein, the integrated circuit and a transfer gate output are connected to data buffer registers. *Id. ¶ 40*, Fig. 10.

2. Overview of Amidi

Amidi discloses a memory interface system with a processor, a memory controller, and a memory module. Ex. 1008 ¶ 3. According to Amidi, a prior art memory interface system is shown in Figure 1, reproduced below.

IPR2014-00883
Patent 8,081,536 B1

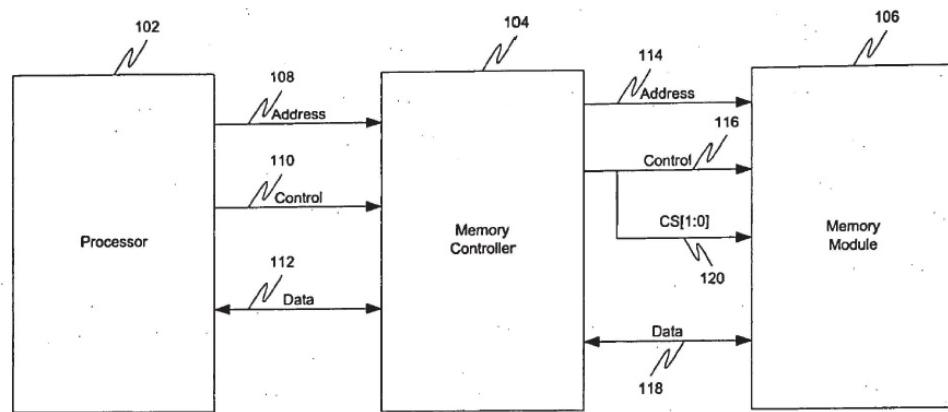


Figure 1 is a schematic of a standard prior art memory interface system.

The prior art system in Figure 1 includes memory module 106 with controller address bus 114, controller control signal bus 116, and controller data bus 118. *Id.* ¶ 3, Fig. 1. As illustrated in Figure 1, memory module 106 communicates with memory controller 104 via busses 114, 116, 118. *Id.* at Fig. 1. Amidi teaches that each stack of DDR memory devices has a data signal line and a data strobe line DQS. *Id.* ¶ 32, Fig. 2. Amidi also teaches that at least two DDR memory devices are connected to a common data memory bus. *Id.* ¶ 34, Fig. 3.

Amidi further discloses multiple memory devices mounted on the front and back side of memory module 400 as shown in Figure 4A reproduced below. *Id.* ¶¶ 34, 37.

IPR2014-00883
Patent 8,081,536 B1

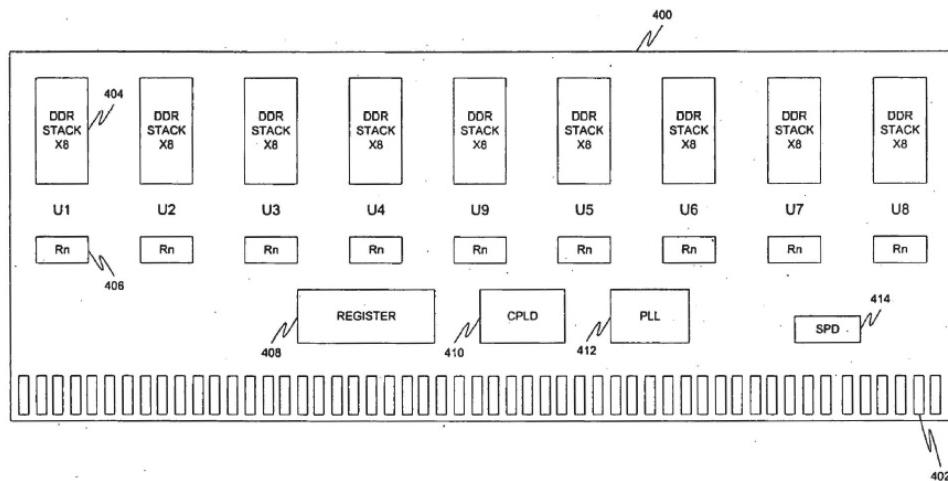


Figure 4A is a schematic of a DDR memory module.

Figure 4A illustrates one embodiment of Amidi where memory module 400 includes memory devices 404, resistor network 406, register 408, complex programmable logic device (CPLD) 410, phase-locked loop (PLL) 412, and SPD 414⁵. *Id.* According to Amidi, memory module 400 receives input signals, including address (Add(n)) signals, row address strobe (RAS) signal, column address strobe (CAS) signal, and bank address (BA[1:0]) signals. Ex. 1008 ¶ 50, Fig. 6A.

Another embodiment of Amidi's memory interface system is shown in Figure 6A, reproduced below.

⁵Amidi discloses that SPD 414 is a simple “I2C interface EEPROM [Electrically Erasable Programmable Read-Only Memory] to hold information regarding memory module for BIOS during the power-up sequence.” Ex. 1008 ¶ 40.

IPR2014-00883
Patent 8,081,536 B1

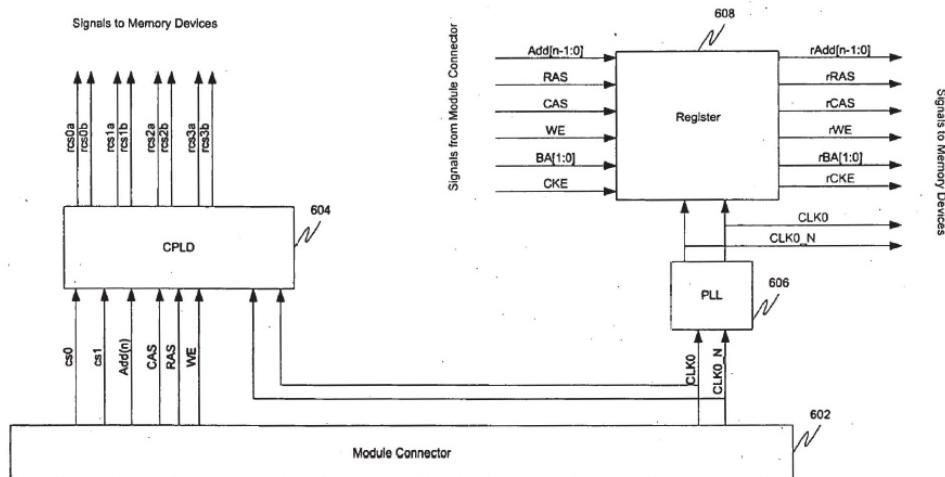


Figure 6A is a schematic of a row address decoding system for a transparent four rank memory module.

As illustrated in Figure 6A, module connector 602 sends signals to CPLD 604, PLL 606, and register 608. *Id.* CPLD 604 also ensures that all commands for a two rank memory module conveyed by module connector 602 are performed on the four rank memory modules. *Id.* ¶ 52. Amidi explains that the system chip select signals control the ranks of individual memory modules. *Id.* ¶¶ 2, 3.

3. Analysis

a. Klein and Amidi Teach or Suggest All the Recited Limitations of Independent Claims 1 and 24

Petitioner contends the combined disclosures of Klein and Amidi, as summarized above, teach or suggest each limitation of independent claims 1, and 24 of the '536 patent. Petitioner first argues that the system described in Klein includes a switch and switch control circuitry interfacing with memory circuits, and, therefore, is a disclosure of a circuit electrically coupled to DDR memory devices. Pet. 18–19 (citing Ex. 1006 ¶ 32). Petitioner then explains that Klein discloses a circuit that is mounted on a memory module and includes a state decoder that may comprise a programmable logic

IPR2014-00883
Patent 8,081,536 B1

device. *Id.* at 13, 22–24; *see Ex.* 1006 ¶ 35). Petitioner also contends that Amidi discloses a circuit mounted on a memory module, where the circuit includes a logic element, a register, and a phase-lock loop device. *Id.* at 22–24 (citing Ex. 1008 ¶¶ 37; 43, 50, 52, Figs. 4A, 6A); Ex. 1011 ¶¶ 58, 64. According to Petitioner, Amidi discloses a memory module having one or more ranks of double-data-rate (DDR) memory devices, which are electrically coupled to the components of the circuit (CPLD). Pet. 23–24 (citing Ex. 1008, ¶¶ 43, 44, 50, 52, Fig. 6A).

Petitioner then argues that Amidi discloses CPLD 604, which activates one of the four ranks with memory devices 306 (“plurality of DDR memory circuits are configured to be activated concurrently with one another for receiving and transmitting data”) in response to four chip-select signals that are generated from two input chip-select signals CS0, CS1, and address signal Add(n). *Id.* at 20, 23–24 (citing Ex. 1008 ¶¶ 43, 52, Fig. 6A). According to Petitioner, Amidi also discloses that PLL generates CLKO and CLKO_N signals (“phase-locked clock signals”) with the chip select signals (rcs0, rcs1, rcs2, rcs3) and relays them to memory devices 306. *Id.* at 21–22, 23–24 (citing Ex. 1008 ¶¶ 50, 52, Fig. 6A). CLKO and CLKO_N signals are also provided to the register and CPLD. *Id.* Petitioner argues that both CPLD 604 and PLL 606 receive clock signals (CLKO and CLKO_N) from module connector 602. *Id.* at 23–24 (citing Ex. 1008, Fig. 6A). Petitioner’s declarant, Dr. Jagannathan, opines that a person of skill in the art would have recognized that the CPLD could receive the clock signals from the PLL. Ex. 1011 ¶ 71.

Finally, Petitioner argues that Klein teaches a circuit that can “selectively isolate” a load because Klein discloses bus switches 32a and

IPR2014-00883
Patent 8,081,536 B1

32b (“the circuit”) that selectively couple data bus segment 30c of memory circuit 28 to input data bus segments 30a, 30b, and decouple (“selectively isolates”) data bus segment 30d of another memory circuit (“a load of the DDR memory circuits of at least one rank”) from the input data bus 30a, 30b that is connected to the memory controller. Pet. 22–23 (citing Ex. 1006 ¶¶ 28, 31, 38, Fig. 6A); Ex. 1011 ¶¶ 72, 98. Petitioner notes Klein specifically states that “[d]ata bus capacitance is reduced by decoupling unaccessed memory circuits from a data bus during data transfers to or from other memory circuits.” Pet. 22 (citing Ex. 1006, Abstract).

According to Petitioner, a person of ordinary skill in the art would have been motivated to combine the teachings of Klein and Amidi, because (1) both references relate to memory devices, (2) both references describe coupling or isolating memory device loads, and (3) the combined teachings would result in the benefit of isolating a memory device load from a computer system so as to reduce parasitic capacitance and increase the speed at which memory accesses can be performed. *Id.* at 26 (citing Ex. 1008 ¶¶ 38, 39, 43, 44, 62; Ex. 1006 ¶¶ 5–10, 28); Ex. 1011 ¶¶ 99, 101.

Petitioner supports its position with the Declaration of Dr. Jagannathan, who testifies that a person of ordinary skill in the art would have had reason to apply the bus switch of Klein to the circuit architecture of Amidi in order to reduce the load seen by the memory controller. Ex. 1011 ¶ 99. Dr. Jagannathan further opines that one of ordinary skill implementing the teachings of Klein would have understood that using a circuit that allows for emulating a higher memory density configuration with lower memory density devices provides the predictable benefit of a cheaper implementation as taught by Amidi. *Id.* ¶ 101.

IPR2014-00883
Patent 8,081,536 B1

Patent Owner contests Petitioner's position, arguing that the combination of Amidi and Klein fails to teach or suggest all the recited claim limitations of claims 1 and 24 and that a person of ordinary skill in the art would not have had a reason to combine the disclosures of the cited references. PO Resp. 25–32.

Patent Owner first contends that Klein fails to teach or suggest “selectively isolate a load of DDR memory circuits,” as recited in the challenged claims. PO Resp. 4–5. Patent Owner argues that different portions of Klein should not be joined together to meet the claim limitations directed to “address signals and a second number of DDR chip-select signals” because a person of ordinary skill in the art would not conflate the embodiments disclosed in Figures 6, 7, and 9 of Klein. *Id.* at 5. According to Patent Owner, Figures 6, 7, and 9 of Klein “are disclosed disparately and their conflation would destroy an intended design purpose of Figure 6 as Klein makes clear.” *Id.* (citing Ex. 1006 ¶¶ 35, 36, 38; Ex. 2002 ¶¶ 98–103). Patent Owner explains that the embodiments illustrated in Figures 5, 6, and 9 of Klein could not be combined without going against the design purpose of Figure 6, which is to avoid an “unconventional” signal line and to not require modification of a DRAM to memory controller interface. *Id.* at 5–11.

Patent Owner relies on the Declaration of Dr. Sechen to support its position. Dr. Sechen testifies that “[a person of ordinary skill in the art] would have understood that the intended purpose of Klein’s Fig. 6 is to promote interoperability by omitting unconventional signals.” Ex. 2002 ¶ 99 (citing Ex. 1006 ¶ 35). According to Dr. Sechen, memory module 60 of Figure 5 has an unconventional signal line because gate control signal 72 is

IPR2014-00883
Patent 8,081,536 B1

received as an input to memory module 60, whereas in Figure 6, gate control signal 72 is generated in memory module 76 by state decoder 78. *Id.* ¶ 100 (citing Ex. 1006, Figs. 5, 6). Dr. Sechen then testifies that in Figure 6, there is no unconventional signal line for gate control (e.g., gate control signal 72) that needs to be created and routed to the memory module. *Id.* (citing Ex. 1006 ¶ 35). Dr. Sechen further states:

Even if one moved the Fig. 9 circuitry of Klein into the Fig. 6 circuitry somehow, Dr. Jagannathan’s position—that Klein discloses selecting a memory module in response to the chip-select signal from the state decoder 78 of Fig. 6 and in response to address bus 87 from the decode circuit 86 of Fig. 9—would lead directly to an inoperable circuit. In this strange case, two sets of TE signals would be generated, one based on control signals and one based on address signals, with no clear way to determine which is the correct TE signal to use for any given set of address and control signals. There is no reasonable way for this to lead to a properly functioning circuit.

Id. ¶ 101. Dr. Sechen, thus, concludes that a person of ordinary skill in the art would not have understood Klein to include selecting a memory module in response to the chip-select signal from state decoder 78 of Figure 6 and in response to address bus 87 from decode circuit 86 of Figure 9, because it would be contrary to Klein’s intended purpose. *Id.* ¶ 102.

We are not persuaded by Patent Owner’s contentions that Klein’s switch is incapable of using address signals *in addition* to chip select signals. We do not find the embodiments illustrated in Figures 4, 5, 6, 7, and 9 to be disparate embodiments that teach away from each other. Rather, we agree with Petitioner’s position that Klein provides a variety of design examples that are intended to be versatile in their application. *See* Pet. 12. In particular, we credit the testimony of Dr. Jagannathan who explains the

IPR2014-00883
Patent 8,081,536 B1

various applications of Klein's teachings and how they would function together. *See e.g.*, Ex. 1026 ¶¶ 19–21 (Supplemental Declaration of Dr. Srinivasan Jagannathan).

Furthermore, we disagree with Patent Owner's contention that the intended design purpose of Klein is to avoid using an unconventional signal line for switch control. *See* PO Resp. 5–11. To the contrary, Klein specifically states that its design purpose is to reduce the parasitic capacitance of a data bus (i.e., capacitive load) (*see* Ex. 1006 ¶ 10), and we credit the testimony of Dr. Jagannathan, who explains that combining the decoder circuit of Klein Figures 6 and 9 does not destroy this stated purpose (Ex. 1026 ¶ 44). Additionally, the description of Figure 6 does not support Patent Owner's position regarding Klein's intended purpose, because Klein teaches merely that “[t]his embodiment has the advantage that no unconventional signal line for gate control needs to be created and routed to the memory module.” *See* Ex. 1001 ¶ 35 (emphasis added), Fig. 6. Noting a design advantage is not the same as indicating a design purpose. *See e.g.*, *In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004) (reference does not teach away if it merely expresses a general preference for an alternative invention, but does not “criticize, discredit, or otherwise discourage” investigation into the invention claimed).

Patent Owner then contends that Amidi fails to teach or suggest “selectively isolate a load of DDR memory circuits,” as recited in the challenged claims. PO Resp. 22–28. Patent Owner specifically argues that Amidi discloses only permanent hard-wiring of all four memory ranks to a data bus and that choosing a rank of memory devices while inactivating other ranks is not selectively isolating a load of DDR memory circuits,

IPR2014-00883
Patent 8,081,536 B1

because neither involves electrical separation from the computer system. *Id.* at 22–23. Patent Owner asserts that direct hard-wiring is permanent, does not allow for separation, and is not subject to acts of selectively isolating (i.e., it does not respond to a selection). *Id.* at 23. Patent Owner, thus, concludes that hard-wiring between the memory ranks and the data bus is a permanent electrical connection that cannot be “selectively isolate[d].” *Id.*

Patent Owner further argues that the proper kind of isolation specified by the challenged claims is load isolation, not merely any kind of isolation, such as some kind of rank-inactivation isolation. *Id.* at 24. According to Patent Owner, a person of ordinary skill in the art would understand that load isolation is directed specifically to the issue of electrical loading. *Id.* (citing Ex. 2002 ¶ 67). Patent Owner concludes that “[d]ue to the permanent hard-wiring of the memory device’s data pin to its line of Amidi’s 72-line data bus, its electrical load is always electrically connected, not separated nor isolated, regardless of any rank activation or inactivation in Amidi. *Id.* at 25.

We do not agree with Patent Owner’s position that hard-wired data signal lines cannot be electrically isolated in a selective fashion. As discussed above, we construe “selectively isolate” as “electrical separation of a selected component from another component.” *See supra*, Section II.A. Amidi’s disclosure of directing electrical signals down a specific signal line or data bus in order to electrically activate a rank within the memory devices and electrically deactivate other ranks falls within the scope of the term “selectively isolate” as we have construed the term. We are further unpersuaded by Patent Owner’s argument that “selectively isolate” is limited narrowly to a particular kind of load isolation. The ’536 patent defines the

IPR2014-00883
Patent 8,081,536 B1

term “load” broadly and includes “*without limitation*, electrical load, such as capacitive load, inductive load, or impedance load.” Ex. 1001, 5:39–41 (emphasis added). Thus, Patent Owner’s arguments are not commensurate with the ’536 patent definition of “load” or with the scope of the challenged claims.

Patent Owner also contends that Klein and Amidi are improperly combined by Petitioner. PO Resp. 28. Patent Owner explains that the Petitioner’s combination is internally inconsistent and based on impermissible hindsight. *Id.* at 28–32. Specifically, Patent Owner argues that the combination of Klein and Amidi is inconsistent, because for a “first number of DDR chip-select signals” the Petitioner relies on off-module signals in Klein and on-module signals in Amidi. *Id.* at 29. According to Patent Owner, it would have been internally inconsistent and impossible to have the “first number of DDR chip-select signals:” be both off-module on-module. *Id.* at 29–30.

Again, we do not agree with Patent Owner. Rather, we agree with Petitioner’s position and we find that Klein’s Figure 10 indicates that the switch and control logic can be incorporated *inside* each memory device, and therefore need not be generated *off*-module. *See* Ex. 1006 ¶¶ 39, 40, Fig. 10; Ex. 1025, 69:13–70:16 (Dr. Sechen’s Deposition Transcript); Reply 16. Klein Figure 10’s teachings demonstrates that the control logic of Figures 7–9 can be combined and even incorporated inside individual memory devices. Patent Owner has not provided sufficient rationale as to why the control signal circuitry shown in Figures 4 and 9 cannot be implemented on module in conformance with Klein’s teachings that “all of the circuitry shown in FIG. 3 may be placed on a single IC, or may be

IPR2014-00883
Patent 8,081,536 B1

provided in a multi-chip package.” Ex. 1006 ¶ 29. Therefore, we are satisfied that use of off-module signals in Klein and on-module signals in Amidi would not be inconsistent with the teachings of Klein.

Finally, Patent Owner contends that a person of ordinary skill in the art would not combine Klein and Amidi because there is a conflict of which “transfer enable” or TE signal to use at Klein’s transfer gates. PO Resp. 31–32 (citing Ex. 2002 ¶¶ 107, 108, 109). Patent Owner relies on the Declaration of Dr. Sechen to support its position regarding the inoperability of a Klein-Amidi combination. *See* Ex. 2002 ¶¶ 107–109. Dr. Sechen specifically states that:

I see that the Patent Owner has identified a technical problem: “a conflict of which TE signal to use at Klein’s transfer gates,” a TE signal from state decoder 78 in Fig. 6 or a TE signal from decode circuit 86 in Fig. 9. (Prelim. Resp., 32.) A [person of ordinary skill in the art] would recognize that such a conflict would not make technical sense, further indicating that the Petitioner and Board are misreading Klein.

Id. ¶ 108. Patent Owner concludes that resolving such a conflict would (i) require undue experimentation by a person of ordinary skill in the art, (ii) change the operating principle of Klein’s transfer gates, and (iii) introduce purposeless redundancy, which would not be obvious to a person of ordinary skill in the art. PO Resp. 32.

We are not persuaded by Patent Owner’s position. Rather, we are satisfied that Klein, specifically, teaches timing control of a transfer gate switch. *See* Ex. 1006 ¶ 23, Fig. 1. We credit the testimony of Dr. Jagannathan, who testifies that Klein teaches how to time the opening and closing of a transfer gate switch. Ex. 1026 ¶¶ 46, 47 (citing Ex. 1006, Fig. 1). Dr. Jagannathan, further testified that “when Klein refers to a ‘signal’

IPR2014-00883
Patent 8,081,536 B1

used in any of the circuits taught therein, one of ordinary skill would understand it is ‘a varying electrical impulse that conveys information from one point to another.’” Ex. 1011 ¶ 89; *see e.g.*, Ex. 1026 ¶¶ 45–47. Thus, we are not persuaded that Amidi’s CS signal and Klein’s TE (“transfer enable”) signal are not equivalent in their signaling function, or that the teachings of Amidi and Klein would not have been combinable to one of ordinary skill in the art, because a person of ordinary skill in the art would have known how to address the timing control for a transfer gate switch based on the teachings of Klein. *See e.g.*, Ex. 1009 ¶ 26 (explaining that transistors are turned on by asserting the gates 15 via an input “transfer enable” signal line labeled TE in Figure 2 and that bus switch circuits such as that illustrated in Figure 2 are known to those of skill in the art).

The test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference. *In re Keller*, 642 F.2d 413, 425 (CCPA 1981). Rather, the test is what the combined teachings of those references would have suggested to those of ordinary skill in the art. *Id.* Patent Owner has not argued that Klein and Amidi teach away from each other. Also, we credit the testimony of Dr. Jagannathan, who states that “[i]t would have been obvious to one of ordinary skill that this teaching of Klein can be applied to Amidi, whereby these elements may be incorporated into a memory integrated circuit.” Ex. 1011 ¶ 99. Therefore, we are persuaded that the teachings of Klein and Amidi are combinable.

Based on the evidence of record, we agree with Petitioner’s position that challenged independent claims 1 and 24 would have been obvious over Klein and Amidi. First, we are persuaded by Petitioner’s reasoning and the

IPR2014-00883
Patent 8,081,536 B1

evidentiary record that Klein discloses a switch and switch control circuitry that are provided on a memory module. We are further persuaded that Amidi teaches a circuit with a logic element and PLL that is mounted on a memory module and that is electrically coupled to four memory ranks made of a plurality of DDR memory devices. We also are persuaded that the teachings of Klein could have been implemented using the circuitry disclosed in Amidi so as to selectively isolate a load of the DDR memory devices from a computer system in response at least in part to a set of signals as required in the challenged claims. Additionally, we credit the testimony of Dr. Jagannathan that a person of ordinary skill in the art would have had a reason to combine the teachings of Klein and Amidi, which both relate to memory devices, such as DIMMS and describe isolating memory device loads. *See* Ex. 1011 ¶¶ 99, 101.

Second, despite Patent Owner's argument to the contrary (*see* PO Resp. 33–35), the arguments presented by Patent Owner generally attack the references individually, rather than in combination. Nonobviousness cannot be established by attacking the references individually when a challenge is predicated upon a combination of prior art disclosures. *See In re Merck & Co., Inc.*, 800 F.2d 1091, 1097 (Fed. Cir. 1986) (citing *In re Keller*, 642 F.2d at 425). In attacking the references individually, Patent Owner fails to address Petitioner's actual challenges and establish an insufficiency in the combined teachings of the references and show Petitioner has not met its burden in arguing obviousness of the challenged claims.

Lastly, we note that the testimony of Patent Owner's Declarant, Dr. Sechen, is based on the claim constructions proffered by Patent Owner and not on the constructions set forth in the Decision to Institute. *Compare* Ex.

IPR2014-00883
Patent 8,081,536 B1

2002 ¶¶ 38–72, with Dec. to Inst. 8–10. We have considered, however, the relevant portions of Dr. Sechen’s analysis regarding the prior art and alleged non-obviousness of the claims and accorded the appropriate weight to that particular testimony.

Accordingly, we hold that Petitioner has shown by a preponderance of the evidence that independent claims 1 and 24 of the ’536 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Klein and Amidi.

b. Klein and Amidi Teach or Suggest All the Recited Limitations of Dependent Claims 16, 17, 30, and 31

Claim 16 depends from claim 1 and recites “[t]he circuit of claim 1, wherein the memory module has attributes.” Ex. 1001, 42:46–51. Petitioner contends the combined disclosures of Klein and Amidi, as summarized above, teach or suggest each limitation of dependent claims 16, 17, 30, and 31 of the ’536 patent. Pet. 34–36. Petitioner specifically contends that Amidi references JEDEC21 C 4-20-4 and Amidi discloses JEDEC compliant DDR DIMMS. *Id.* at 34. According to Petitioner JEDEC21C 4-20-4 defines the specifications for DDR DIMMs, and discloses attributes associated with the plurality of DDR memory devices. *Id.* Thus, Petitioner concludes that Amidi in view of JEDEC21 C 4-20-4 teaches the further limitations of dependent claim 16. *Id.*

Claim 17 depends from claim 16 and recites a circuit of claim 16, wherein the attributes are selected from a group consisting of: a number of row addresses, a number of column addresses, a number of DDR memory circuits, a data width of the DDR memory circuits, a memory density per DDR memory circuit, a number of ranks, and a memory density per

IPR2014-00883
Patent 8,081,536 B1

rank. Ex. 1001, 42:52–57. Petitioner contends Amidi and Klein meet the limitations of claim 17, because Amidi teaches the following:

the memory module has at least the following attributes: a number of row addresses, a number of column addresses, thirty-six memory devices (“a number of DDR memory circuits”), 8 data bits (“a data width per DDR memory circuit”), 16 M Byte memory density per memory device (“a memory density per DDR memory circuit”), four ranks (rank 0, rank 1, rank 2, rank 3) (“a number of ranks of DDR memory circuit”),

Id. at 36 (citing Ex. 1008 ¶ 8). According to Petitioner, it would have been obvious to one of ordinary skill in the art that the attributes include at least the number of row address signals. *Id.* (citing Ex. 1001 ¶ 57).

Claims 30 and 31 are method claims that are substantially the same as claims 16 and 17. Ex. 1001, 43:63–44:6.

Patent Owner does not provide separate contentions regarding additional limitations recited in the dependent claims. *See generally* PO Resp.

After consideration of the language recited in claims 16, 17, 30, and 31 of the ’536 patent, the Petition, the Patent Owner Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers, we find that one of ordinary skill in the art would have considered these dependent claims obvious over the combination of Klein and Amidi. Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that dependent claims 16, 17, 30, and 31 of the ’536 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Klein and Amidi.

IPR2014-00883
Patent 8,081,536 B1

F. Asserted Obviousness of Claims 16, 17, 30, and 31 in view of Klein, Amidi, and Dell

Petitioner contends claims 16, 17, 30, and 31 of the '536 patent are unpatentable under 35 U.S.C. § 103 in view of Klein, Amidi, and Dell. Pet. 27–33. Patent Owner disputes Petitioner’s position, arguing that Dell fails to make up for the deficiencies of Klein and Amidi with respect to certain claim limitations. PO Resp. 35–36.

We have reviewed the Petition, the Patent Owner’s Response, and Petitioner’s Reply, as well as the relevant evidence discussed in those papers. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that claims 16, 17, 30, and 31 of the '536 patent are unpatentable as obvious over the combination of Klein, Amidi, and Dell.

1. Overview of Klein

See Section II.D.1., discussed above.

2. Overview of Amidi

See Section II.D.2., discussed above.

3. Overview of Dell

Dell discloses a system for address re-mapping for memory modules using presence detect data. Ex. 1009, Abstract. Dell discloses a memory module having a presence detect (PD) that allows a four bank memory device to be used with a computer system that is expecting a two bank memory device. *Id.* at 3:30–33; 4:55–58. The PD data provides the operational characteristics and compatibility with system requirements of the memory module. *Id.* at 2:8–12. Dell discloses that memory module 20, shown in Figure 1 reproduced below, includes logic circuit 24 that interfaces

IPR2014-00883
Patent 8,081,536 B1

with memory controller 14 of CPU system 12 using address, data, and bus control signals. *Id.* at 3:66–4:5, 3:40–42.

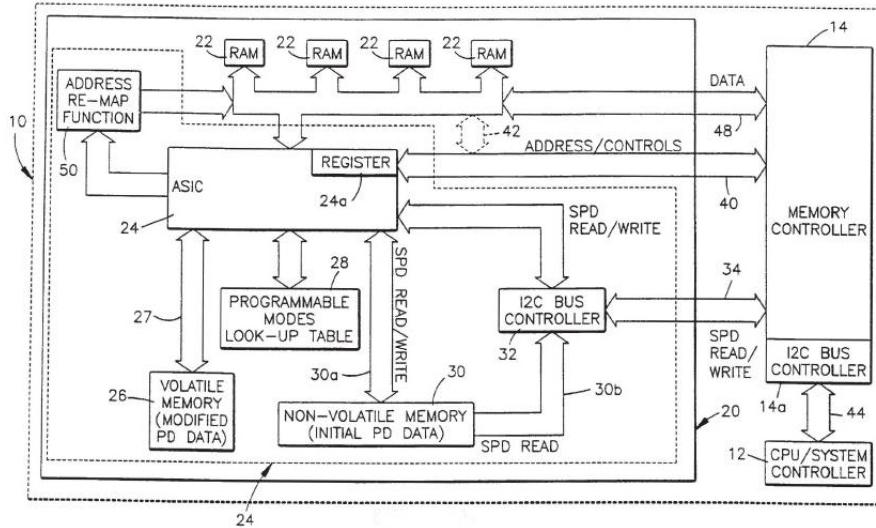


Figure 1 is a schematic of memory module 20 connected to computer system 12.

As illustrated in Figure 1, memory module 20 includes a plurality of memory devices being configured as memory banks 22. *Id.* at 2:48–51. The memory devices may be DDR synchronous DRAM (SDRAM). *Id.* at 1:56–67. Dell discloses that memory module 20 includes non-volatile memory 30 for storing serial PD (SPD) data that is accessible to system controller 12. *Id.* at 5:28–32. Memory module 20 includes volatile memory 26 for storing modified PD data that is transferred from CPU system 12 based on system requirements. *Id.* at 7:5–16.

4. Analysis

Claim 16 depends from claim 1 and recites “[t]he circuit of claim 1, wherein the memory module has attributes.” Ex. 1001, 42:46–51. Petitioner contends that Dell discloses a memory module with attributes, because the memory module of Dell has (1) presence detect (PD) data, and (2) memory device characteristics determined from serial presence detect (SPD) data.

IPR2014-00883
Patent 8,081,536 B1

Pet. 27. Claim 16 further requires “the circuit in the at least one configuration is further configured to store data accessible to the computer system.” Petitioner contends that Dell discloses that non-volatile memory 30 stores serial presence detect (SPD) data that is accessible to system controller 12, thereby meeting the required elements of claim 16. *Id.* at 28.

Claim 17 depends from claim 16 and recites a circuit of claim 16, wherein the attributes are selected from a group consisting of: a number of row addresses, a number of column addresses, a number of DDR memory circuits, a data width of the DDR memory circuits, a memory density per DDR memory circuit, a number of ranks, and a memory density per rank. Ex. 1001, 42:52–57. Petitioner contends Dell meets the limitations of claim 17, because Dell discloses that a SDRAM 22 uses twelve address signals A0-A11 and Dell discloses storing modified PD data as the up-to-date PD data and performing an address-remapping function 50. Pet. 31.

Claims 30 and 31 are method claims that are substantially the same as claims 16 and 17. Ex. 1001, 43:63–44:6.

According to Petitioner, a person of ordinary skill in the art would have been motivated to combine the teachings of Klein, Amidi, and Dell, because (1) all three references relate to memory devices, and (2) a person of skill in the art would have been motivated to use the non-volatile memory and volatile memory in Dell with the SPD in Amidi and the switches of Klein to have a higher density memory device using a number of lower density memory devices that are comparatively inexpensive and readily available. *Id.* at 32–33 (citing Ex. 1008 ¶ 8; Ex. 1009, 4:55–58); Ex. 1011 ¶¶ 99, 101.

IPR2014-00883
Patent 8,081,536 B1

Petitioner supports its position with the Declaration of Dr. Jagannathan, who testifies that using lower memory density devices poses the challenge of reporting a different configuration to the memory controller. Ex. 1011 ¶ 102. According to Dr. Jagannathan, Dell teaches a solution to that problem. *Id.* Thus, Dr. Jagannathan concludes that one of ordinary skill would seek to apply Dell to Amidi to further improve an actual memory system for emulating a different configuration expected by a host system. *Id.*

Patent Owner contends that Dell fails to cure the deficiencies Patent Owner perceives in Klein and Amidi. PO Resp. 35–36. Therefore, Patent Owner concludes that dependent claims 16, 17, 30, and 31 would not have been rendered obvious for the same reasons cited by Patent Owner in its arguments against the combination of Klein and Amidi. *Id.*

After careful consideration of the language recited in claims 16, 17, 30, and 31 of the '536 patent, the Petition, the Patent Owner Response, and Petitioner's Reply, as well as the relevant evidence discussed in those papers, we find that for the same reasons discussed in detail previously (*see supra*, Section II.D.3), one of ordinary skill in the art would have considered these dependent claims obvious over the combination of Klein, Amidi, and Dell. Accordingly, we determine that Petitioner has shown by a preponderance of the evidence that dependent claims 16, 17, 30, and 31 of the '536 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of Klein, Amidi, and Dell.

IPR2014-00883
Patent 8,081,536 B1

III. CONCLUSION

We conclude Petitioner has shown by a preponderance of the evidence that claims 1, 16, 17, 24, 30, and 31 of the '536 patent are unpatentable under 35 U.S.C. § 103(a) as having been obvious over the combination of (1) Klein and Amidi and (2) Klein, Amidi, and Dell.

IV. ORDER

For the reasons given, it is

ORDERED that, by a preponderance of the evidence, claims 1, 16, 17, 24, 30, and 31 of the '536 patent are unpatentable; and

FURTHER ORDERED that because this is a Final Written Decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2014-00883
Patent 8,081,536 B1

FOR PETITIONER:

Fabio Marino
fmarino@mwe.com

Bernie Knight
bknights@mwe.com

Barrington Dyer
bdyer@mwe.com

MCDERMOTT WILL & EMERY

FOR PATENT OWNER:

Mehran Arjomand
marjomand@mofo.com
patentdocket@mofo.com

David Kim
dkim@mofo.com
patentdocket@mofo.com

Erol Basol
ebasol@mofo.com
patentdocket@mofo.com

Jean Nguyen
jnguyen@mofo.com
patentdocket@mofo.com

Jonathan Statman
jstatman@mofo.com
patentdocket@mofo.com

MORRISON & FOERSTER LLP



US007881150B2

**(12) United States Patent
Solomon et al.**

**(10) Patent No.: US 7,881,150 B2
(45) Date of Patent: *Feb. 1, 2011**

(54) CIRCUIT PROVIDING LOAD ISOLATION AND MEMORY DOMAIN TRANSLATION FOR MEMORY MODULE

(75) Inventors: Jeffrey C. Solomon, Irvine, CA (US); Jayesh R. Bhakta, Cerritos, CA (US)

(73) Assignee: Netlist, Inc., Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: 12/629,827

(22) Filed: Dec. 2, 2009

(65) Prior Publication Data

US 2010/0128507 A1 May 27, 2010

Related U.S. Application Data

(63) Continuation of application No. 12/408,652, filed on Mar. 20, 2009, now Pat. No. 7,636,274, which is a continuation of application No. 11/335,875, filed on Jan. 19, 2006, now Pat. No. 7,532,537, which is a continuation-in-part of application No. 11/173,175, filed on Jul. 1, 2005, now Pat. No. 7,289,386, said application No. 11/173,175 is a continuation-in-part of application No. 11/075,395, filed on Mar. 7, 2005, now Pat. No. 7,286,436.

(60) Provisional application No. 60/645,087, filed on Jan. 19, 2005, provisional application No. 60/588,244, filed on Jul. 15, 2004, provisional application No. 60/550,668, filed on Mar. 5, 2004, provisional application No. 60/575,595, filed on May 28, 2004, provisional application No. 60/590,038, filed on Jul. 21, 2004.

**(51) Int. Cl.
G11C 8/16** (2006.01)

**(52) U.S. Cl. 365/233.13; 365/51; 365/149;
365/230.03; 365/230.06**

**(58) Field of Classification Search 365/51,
365/149, 230.03, 230.06, 233.13**
See application file for complete search history.

(56) References Cited

U.S. PATENT DOCUMENTS

4,368,515 A 1/1983 Nielsen

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 92/02879 2/1992

(Continued)

OTHER PUBLICATIONS

Karabatsos, C., "Quad Band Memory (QBM) Technology," Kentron Technologies, Inc., Apr. 2001, pp. 1-5.*

(Continued)

Primary Examiner—Dang T Nguyen

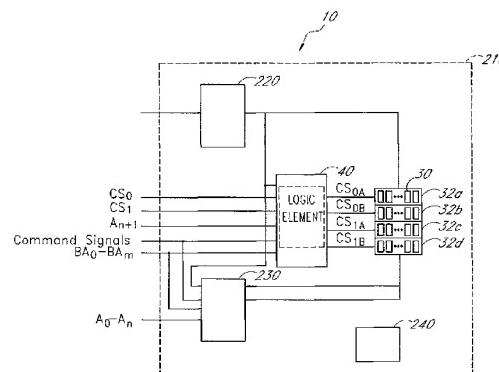
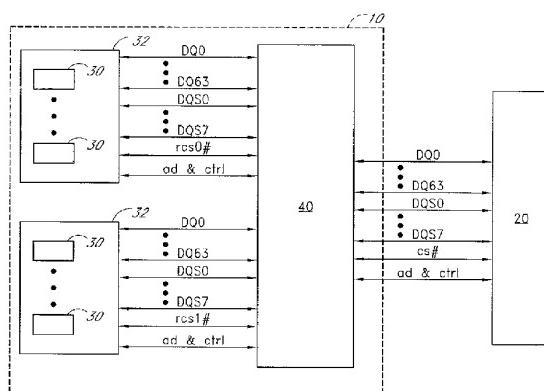
Assistant Examiner—Alexander Sofocleous

(74) Attorney, Agent, or Firm: Knobbe, Martens, Olson & Bear, LLP

(57) ABSTRACT

A circuit is configured to be mounted on a memory module so as to be electrically coupled to a plurality of double-data-rate (DDR) memory devices arranged in one or more ranks on the memory module. The circuit includes a logic element, a register, and a phase-lock loop device. The circuit is configurable to respond to a set of input signals from a computer system to selectively isolate one or more loads of the plurality of DDR memory devices from the computer system and to translate between a system memory domain of the computer system and a physical memory domain of the plurality of DDR memory devices.

36 Claims, 23 Drawing Sheets



US 7,881,150 B2

Page 2

U.S. PATENT DOCUMENTS					
4,392,212 A	7/1983	Miyasaka et al.	6,742,098 B1	5/2004	Halbert et al.
4,633,429 A	12/1986	Lewandowski et al.	6,754,797 B2	6/2004	Wu et al.
4,670,748 A	6/1987	Williams	6,785,189 B2	8/2004	Jacobs et al.
4,866,603 A	9/1989	Chiba	6,788,592 B2	9/2004	Nakata et al.
4,958,322 A	9/1990	Kosugi et al.	6,807,125 B2	10/2004	Coteus et al.
4,961,172 A	10/1990	Shubat et al.	6,813,196 B2	11/2004	Park et al.
4,980,850 A	12/1990	Morgan	6,834,014 B2	12/2004	Yoo et al.
5,247,643 A	9/1993	Shottan	7,356,639 B2	12/2004	Perego et al.
5,345,412 A	9/1994	Shiratsuchi	6,854,042 B1	2/2005	Karabatosos
5,426,753 A	6/1995	Moon	6,880,094 B2	4/2005	LaBerge
5,483,497 A	1/1996	Mochizuki et al.	6,889,304 B2	5/2005	Perego et al.
5,495,435 A	2/1996	Sugahara	6,912,615 B2	6/2005	Nicolai
5,581,498 A	12/1996	Ludwig et al.	6,912,628 B2	6/2005	Wicki et al.
5,590,071 A	12/1996	Kolor et al.	6,925,028 B2	8/2005	Hosokawa et al.
5,699,542 A	12/1997	Mehta et al.	6,944,694 B2	9/2005	Pax
5,702,984 A	12/1997	Bertin et al.	6,950,366 B1	9/2005	Lapidus et al.
5,703,826 A	12/1997	Hush et al.	6,961,281 B2	11/2005	Wong et al.
5,745,914 A	4/1998	Connolly et al.	6,981,089 B2	12/2005	Dodd et al.
5,802,395 A	9/1998	Connolly et al.	6,982,892 B2	1/2006	Lee et al.
5,805,520 A	9/1998	Anglada et al.	6,982,893 B2	1/2006	Jakobs
5,822,251 A	10/1998	Bruce et al.	6,990,043 B2	1/2006	Kuroda et al.
RE36,229 E	6/1999	Kang	6,996,686 B2	2/2006	Doblar et al.
5,926,827 A	7/1999	Dell et al.	7,007,130 B1	2/2006	Verbrugge
5,959,930 A	9/1999	Sakurai	7,007,175 B2	2/2006	Chang et al.
5,963,464 A	10/1999	Dell et al.	7,046,538 B2 *	5/2006	Kinsley et al. 365/52
5,966,736 A	10/1999	Gittinger et al.	7,054,179 B2	5/2006	Cogdill et al.
6,018,787 A	1/2000	Ip	7,065,626 B2	6/2006	Schumacher et al.
6,044,032 A	3/2000	Li	7,073,041 B2	7/2006	Dwyer et al.
6,070,217 A	5/2000	Connolly et al.	7,120,727 B2	10/2006	Lee et al.
6,070,227 A	5/2000	Rokicki	7,124,260 B2	10/2006	LaBerge et al.
6,097,652 A	8/2000	Roh	7,127,584 B1	10/2006	Thompson et al.
6,108,745 A	8/2000	Gupta et al.	7,130,952 B2	10/2006	Nanki et al.
6,134,638 A	10/2000	Olarig et al.	7,133,960 B1	11/2006	Thompson et al.
6,151,271 A	11/2000	Lee	7,133,972 B2	11/2006	Jeddeloh
6,154,418 A	11/2000	Li	7,142,461 B2	11/2006	Janzen
6,154,419 A	11/2000	Shakkarwar	7,149,841 B2	12/2006	LaBerge
6,185,654 B1	2/2001	Van Doren	7,167,967 B2	1/2007	Bungo et al.
6,209,074 B1	3/2001	Dell et al.	7,181,591 B2	2/2007	Tsai
6,226,709 B1	5/2001	Goodwin et al.	7,200,021 B2	4/2007	Raghuram
6,226,736 B1	5/2001	Niot	7,227,910 B2	6/2007	Lipka
6,233,650 B1 *	5/2001	Johnson et al. 711/5	7,266,639 B2	9/2007	Raghuram
6,247,088 B1	6/2001	Seo et al.	7,272,709 B2	9/2007	Zitlaw et al.
6,317,352 B1	11/2001	Halbert et al.	7,281,079 B2	10/2007	Bains et al.
6,400,637 B1	6/2002	Akamatsu et al.	7,286,436 B2	10/2007	Bhakta et al.
6,408,356 B1	6/2002	Dell	7,289,386 B2	10/2007	Bhakta et al.
6,414,868 B1	7/2002	Wong et al.	7,346,750 B2	3/2008	Ishikawa
6,415,374 B1	7/2002	Fauve et al.	7,370,238 B2	5/2008	Billick et al.
6,446,158 B1 *	9/2002	Karabatosos 711/5	7,437,591 B1	10/2008	Wong
6,446,184 B2 *	9/2002	Dell et al. 711/170	7,461,182 B2	12/2008	Fukushima et al.
6,453,381 B1	9/2002	Yuan et al.	7,471,538 B2	12/2008	Hofsta
6,470,417 B1	10/2002	Kolor et al.	7,532,537 B2	5/2009	Solomon et al.
6,502,161 B1	12/2002	Perego et al.	7,619,912 B2	11/2009	Bhakta et al.
6,518,794 B2	2/2003	Coteus et al.	7,636,274 B2	12/2009	Solomon et al.
6,526,473 B1	2/2003	Kim	2001/0003198 A1	6/2001	Wu
6,530,007 B2	3/2003	Olarig et al.	2001/0052057 A1	12/2001	Lai
6,530,033 B1	3/2003	Rayham et al.	2002/0088633 A1	7/2002	Kong et al.
6,553,450 B1	4/2003	Dodd et al.	2003/0063514 A1	4/2003	Fauve
6,618,320 B2	9/2003	Hasegawa et al.	2003/0090359 A1	5/2003	OK
6,621,496 B1	9/2003	Ryan	2003/0090879 A1	5/2003	Bhakta et al.
6,625,081 B2	9/2003	Rochparvar et al.	2003/0191995 A1	10/2003	Bhakta et al.
6,625,687 B1	9/2003	Halbert et al.	2003/0210575 A1	11/2003	Seo et al.
6,636,935 B1	10/2003	Ware et al.	2004/0000708 A1	1/2004	Rapport et al.
6,646,949 B1	11/2003	Ellis et al.	2004/0037158 A1	2/2004	Coteus et al.
6,658,509 B1	12/2003	Bonella et al.	2004/0201968 A1	10/2004	Tafolla
6,674,684 B1	1/2004	Shen	2005/0036378 A1	2/2005	Bhakta et al.
6,681,301 B1	1/2004	Mehtha et al.	2005/0281096 A1	12/2005	Bhakta et al.
6,683,372 B1	1/2004	Wong et al.	2006/0044860 A1	3/2006	Kinsley et al.
6,697,888 B1	2/2004	Halbert et al.	2006/0117152 A1	6/2006	Amidi et al.
6,705,877 B1	3/2004	Li et al.	2006/0126369 A1	6/2006	Raghuram
6,717,855 B2	4/2004	Lai	2006/0129755 A1	6/2006	Raghuram
6,738,880 B2	5/2004	Lai et al.	2006/0179206 A1	8/2006	Brittain et al.
			2006/0259711 A1	11/2006	Oh
			2006/0267172 A1	11/2006	Nguyen et al.

US 7,881,150 B2

Page 3

2006/0277355 A1 12/2006 Ellsberry et al.
 2010/0091540 A1* 4/2010 Bhakta et al. 365/51

FOREIGN PATENT DOCUMENTS

WO	WO 94/07242	3/1994
WO	WO 95/34030	12/1995
WO	WO 02/58069	7/2002
WO	WO 03/17283	2/2003
WO	WO 03/69484	8/2003
WO	WO 2006/055497	5/2006

OTHER PUBLICATIONS

- “64 & 72 Pin Zip/Simm Sram Module,” JEDEC, Standard No. 21-C, www.jedec.com/download/search/404_01.pdf, Jun. 1997 pp. 4.4.1-1.
- Abali, B. “Memory Expansion Technology (MXT): Software Support and Performance,” IBM J. Res. & Dev., vol. 45, No. 2, 2001, pp. 287-300.
- Arlington, DL Evans. “Enhancement of Memory Card Redundant Bit Usage Via Simplified Fault Alignment Exclusion,” IBM Technical Disclosure Bulletin, 1987.
- Arroyo et al. “Method of executing Manufacturing ROM Code Without Removing System Roms,” IP.com, IPCOM000037214D, 1989.
- “Bank Striping of Data Across Internal SDRAM Banks,” IP.com, IPCOM000013697D, 2000.
- Barr, Michael. “Programmable Logic: What’s it to Ya?,” Embedded Systems Programming, Jun. 1999, pp. 75-84.
- Bennayoun et al. “Input/Output Chip Select Doubler,” IBM Technical Disclosure Bulletin, vol. 38, No. 04 1995, pp. 237-240.
- Blum et al. “Fast Multichip Memory System With Power Select Signal,” IMB Technical Disclosure Bulletin, 1979.
- Cuppu et al. “A Performance Comparison of Contemporary DRAM Architectures,” IEEE Proceedings of the 26th International Symposium on Computer Architectures, May 2-4, 1999, Atlanta, Georgia, pp. 1-12.
- Cuppu et al. “Concurrency, Latency, or System Overhead: Which Has the Largest Impact on Uniprocessor DRAM-System Performance?,” IEEE, 2001, pp. 62-71.
- Cuppu et al. “High-Performance DRAMs in Workstation Environments,” IEEE Transactions on Computers, vol. 50, No. 11, 2001, pp. 1133-1153.
- Denneau, M. “Logic Processor for Logic Simulation Machine,” IBM Technical Disclosure Bulletin, vol. 25, No. 1, 1982.
- “Distributed Memory Mapping,” IP.com, IPCOM000014788D, 2000.
- Fairchild Semiconductor. “DM74LS138 DM74LS139 Decoder/Demultiplexer,” Fairchild Semiconductor Corporation, 2000.
- Fitzgerald et al. “Chip Select Circuit for Multi-Chip RAM Modules,” IP.com, IPCOM000044404D, 1984.
- Gray, KS. “Fet Ram Chip Double Density Scheme,” IP.com, IPCOM0000043942D, 1984.
- Grimes et al. “Access Rate/Availability Improvement Logic for Dynamic Memories,” IBM Technical Disclosure Bulletin, Oct. 1982.
- Gupta et al. “Designing and Implementing a Fast Crossbar Scheduler,” IEEE Micro, 1999, pp. 20-28.
- Hession et al. “Chip Select Technique for Multi Chip Decoding,” IP.com, IPCOM000070404D, 1985.
- Hewlett-Packard. “Memory technology evolution: an overview of system memory technologies,” technology brief, 7th edition, 2003.
- Hoare et al. “An 88-Way Multiprocessor Within an FPGA With Customizable Instructions,” Proceedings of the 18th International Parallel and Distributed Processing Symposium, 2004.
- “Information Huawei or FPGA-Take Five,” Electronic News, 2002, p. 24.
- Intel Corporation, PC SDRAM Registered DIMM Design Support Document, Revision 1.2, Oct. 1998.
- Intel Corporation, 66/100 MHz PC SDRAM 64-Bit Non-ECC/Parity 144 Pin Unbuffered SO-DIMM Specification, Revision 1.0, Feb. 1999.
- JEDEC Standard No. 21-C, 4.20-2—168 Pin, PC133 SDRAM Registered Design Specification, Revision 1.4, Release 11a, Feb. 2002.
- JEDEC Standard No. 21-C, 4.20-3—144 Pin, PC133 SDRAM Unbuffered SO-DIMM, Reference Design Specification, Revision 1.02, Release 11. Published Oct. 2003.
- JEDEC Standard No. 21-C, DDR SDRAM PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification, Revision 1.3, Release 11b, Jan. 2002.
- JEDEC Standard No. 21-C, 4.20.5-184 Pin, PC1600/2100 DDR SDRAM Unbuffered DIMM Design Specification, Revision 1.1, Release 11b. Published Apr. 2003.
- JEDEC Standard No. 21-C, 4.20.5-184 Pin, PC2700/PC2100/PC1600 DDR SDRAM Unbuffered SO-DIMM Reference Design Specification, Revision 1.1, Release 11 b, Apr. 26, 2002.
- JEDEC Standard JESD79D, “Double Data Rate (DDR) SDRAM Specification,” published Feb. 2004.
- Jin et al. “Embedded Memory in System-On-Chip Design: Architecture and Prototype Implementation,” CCECE, 2003, pp. 141-146.
- Jin et al. “Prototype Implementation and Evaluation of a Multibank Embedded Memory Architecture in Programmable Logic,” IEEE, 2003, pp. 13-16.
- Kane et al. “Read Only Store Memory Extension,” IP.com, IPCOM000082845D, 1975.
- Kornaros et al. “A Fully-Programmable Memory Management System Optimizing Queue Handling at Multi Gigabit Rates,” DAC, 2003, pp. 54-69.
- Lee et al. “A banked-promotion translation lookaside buffer system,” Journal of Systems Architecture, vol. 47, 2002, pp. 1065-1078.
- Lee et al. “An on-chip cache compression technique to reduce decompression overhead and design complexity,” Journal of Systems Architecture, vol. 46, 2000, pp. 1365-1382.
- Lin et al. “Designing a Modern Memory Hierarchy with Hardware Prefetching,” IEEE Transactions on Computers, vol. 50, No. 11, 2001, pp. 1202-1217.
- Luthra et al. “Interface Synthesis Using Memory Mapping for an FPGA Platform,” Proceedings of the 21st International Conference on Computer Design, 2003.
- Matick et al. “Read-Select Capability for Static Random-Access Memory,” IMB Technical Disclosure Bulletin, 1985, pp. 6640-6642.
- Matick, RE. “Logic and Decoder Arrangement for Controlling Spill/Wrap Boundaries of a Bit-Addressable Memory Decoder,” IMB Technical Disclosure Bulletin, 1984.
- “Method for a high-performance DRAM address mapping mechanism,” IP.com, IPCOM000008164D, 2002.
- “Method for memory probing on a multiple-DIMM bus,” IP.com, IPCOM000019063D, 2003.
- “Method for multiple device interface testing using a single device,” IP.com, IPCOM000010054D, 2002.
- Meyers et al. “Use of Partially Good Memory Chips,” IP.com, IPCOM000066246D, 1979.
- Murdocka et al. “Principles of Computer Architecture,” Prentice Hall, 2000, pp. 249-251.
- Ofek et al. “Partial Two Way Mapping Technique,” IMB Technical Disclosure Bulletin, 1969.
- Paldan, David. “Programmable Memory Address Decoding For Microprocessor Memory Device,” IP.com, IPCOM000005486D, 1983.
- “PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification” JEDEC, Standard No. 21-C, Revision 1-3, Jan. 2002, pp. 4.20.4-1.
- Pellinger et al. “Dual Addressable Memory,” IP.com, IPCOM000068610D, 1978.
- Plotnick et al. “Shuffle Your Chips For Better Performance,” PC Week, 1998, p. 90.
- Schubert et al. “Accelerating system integration by enhancing hardware, firmware, and co-simulation,” IBM J. Res. & Dev, vol. 48, No. 3/4, May/Jul. 2004, pp. 569-581.
- Skelton, MH. “Program Controlled Paging Scheme for Memory Expansion,” IP.com, IPCOM000050954D, 1982.
- Slegel et al. “IBM’s S/390 G5 Microprocessor Design,” IEEE Micro, 1999, pp. 12-23.
- Smith, BA. “Chip Select Decoder Circuit,” IP.com, IPCOM000063400D, 1985.
- Stelzer, KC. “Planar Memory Boundary Registers with Remap Feature,” IMB Technical Disclosure Bulletin, 1993.

US 7,881,150 B2

Page 4

- Sunaga et al. "An Enable Signal Circuit for Multiple Small Banks," IP.com, IPCOM000015887D, 2002.
- Sunaga et al. "Continuous RAS Access Method in Multiple-bank DRAM Chip," IP.com, IPCOM000123375D, 1998.
- Toal et al. "A 32-Bit SoPC Implementation of a P5." Proceedings of the Eighth IEEE International Symposium on Computers and Communications, 2003, pp. 1530-1546.
- Tudruj, Marek. "Dynamically reconfigurable heterogeneous multi-processor systems with transputer-controlled communication." Journal of Systems Architecture, vol. 43, 1997, pp. 27-32.
- Yao, YL. High Density Memory Selection Circuit,: IP.com, IPCOM000078218D, 1972.
- Google, Inc. v. Netlist, Inc.*, No. 4:08-cv-04144-SBA. Netlist Inc.'s Answer to Complaint and Counterclaim (N.D. Ca. Filed Nov. 18, 2008).
- Google, Inc. v. Netlist, Inc.*, No. C 08-04144 SBA Google Inc.'s Invalidity Contentions Pursuant to PAT. L.F. 3-3, dated Apr. 13, 2009.
- Google, Inc. v. Netlist, Inc.*, No. C08 04144, Complaint for Declaratory Relief, (N.D. Ca Dated Aug. 29, 2008).
- Letter from G. Hopkins Guy III, Orrick, Herrington & Sutcliffe LLP, to R. Scott Oliver, Morrison & Foerster, (Apr. 14, 2009).
- MetaRAM, Inc. v. Netlist, Inc.* No. 3:09-cv-01309-VRW, MetaRAM's Reply to Netlist's Counterclaims, (N.D. Ca. Filed Jun. 3, 2009).
- MetaRAM, Inc. v. Netlist, Inc.*, No. 3:09-cv-01309-VRW, Netlist's Answer to Complaint and Counterclaims, (N.D. Ca, filed May 11, 2009).
- MetaRAM, Inc. v. Netlist, Inc.*, No. C09 01309, Complaint for Patent Infringement, (N.D. Ca. Filed Mar. 25, 2009).
- Netlist, Inc. v. MetaRAM, Inc.*, No. 09-165-GMS, MetaRAM, Inc.'s Answer and Affirmative Defenses to Plaintiff's Complaint, dated Apr. 20, 2009.
- Netlist, Inc. v. MetaRAM, Inc.*, No. 1:09-cv-00165-GMS, Complaint for Patent Infringement, (D. Del. Filed Mar. 12, 2009).
- U.S. Appl. No. 12/504,131; filed Jul. 16, 2009; Owned by Netlist, Inc.
- U.S. Appl. No. 12/761,179; filed Apr. 15, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 12/815,339; filed Jun. 14, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 12/774,632; filed May 5, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 12/422,925; filed Apr. 13, 2009; Owned by Netlist, Inc.
- U.S. Appl. No. 12/422,853; filed Apr. 13, 2009; Owned by Netlist, Inc.
- U.S. Appl. No. 95/001,339; filed Jun. 8, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 95/001,337; filed Jun. 4, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 95/000,546; filed May 11, 2010; Owned by Netlist, Inc.
- U.S. Appl. No. 95/001,381; filed Jun. 9, 2010; Owned by Netlist, Inc.
- "Quad Band Memory (QBMMAT™): DDR200/266/333 devices producing DDR400/533/667" (the "QBMA Reference"), published by the QBMA Alliance, Platform Conference, San Jose, California, Jan. 23-24, 2002.
- Carvalho, Carlos; "The Gap between Processor and Memory Speeds"; ICCA '02.
- Kirihata et al.; "A 390-mm, 16-Bank, 1-Gb DDR SDRAM with Hybrid Bitline Architecture"; IEEE Journal of Solid-State Circuits, vol. 34, No. 11, Nov. 1999.
- Jacob, Bruce L.; "Synchronous DRAM Architectures, Organizations, and Alternative Technologies". University of Maryland, Dec. 10, 2002.
- Keltcher et al.; "The AMD Opteron Processor for Multiprocessor Servers"; IEEE Computer Society.
- Kellogg, Mark; "PC133: SDRAM Main Memory Performance Reaches New Heights"; IBM Microelectronics, 1999.
- Jedec "JEDEC Standard: Double Data Rate (DDR) SDRAM Specification", JESD79C Mar. 2003.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*. Complaint for Patent Infringement, filed Sep. 22, 2009 in 10 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*, Defendant Inphi Corporation's Answer To Plaintiffs Complaint for Patent Infringement, filed Nov. 12, 2009 in 6 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*, Plaintiff Netlist, Inc's First Amended Complaint For Patent Infringement, filed Dec. 23, 2009 in 8 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*, Defendant Inphi Corporation's Answer To Plaintiffs First Amended Complaint For Patent Infringement, filed Feb. 11, 2010 in 9 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*, Defendant Inphi Corporation 's Notice Of Motion and Motion For Stay Pending Reexaminations and Interference Proceeding Regarding The Patents-In-Suit; Memorandum Of Points and Authorities In Support Thereof, filed Apr. 21, 2010 in 28 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *NETLIST, INC. vs. INPHI CORPORATION*, Plaintiff Netlist Inc's Opposition To Defendant Inphi Corporation's Motion For Stay Pending Reexaminations And Interference Proceedings Regarding The Patents-In-Suit, filed May 3, 2010 in 23 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Complaint For Patent Infringement, filed Dec. 4, 2009 in 47 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Google's Answer To Plaintiffs Complaint For Patent Infringement; And Assertion of Counterclaims, filed Feb. 12, 2010 in 13 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Plaintiff Netlist, Inc.'s Reply To Defendant Google Inc.'s Counterclaim, filed Mar. 8, 2010 in 11 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Joint Claim Construction And Prehearing Statement Under Patent Local Rule 4-3, filed Jun. 25, 2010 in 5 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Exhibit A to Joint Claim Construction and Prehearing Statement under Patent L.R. 4-3, filed Jun. 25, 2010 in 2 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Exhibit B to Joint Claim Construction and Prehearing Statement under Patent L.R. 4-3, filed Jun. 25, 2010 in 23 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Plaintiff Netlist, Inc's Opening Claim Construction Brief, filed Jul. 16, 2010 in 29 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Defendant Google Inc's Responsive Claim Construction Brief, filed Aug. 4, 2010 in 27 pages.
- U.S. District Court Northern District of California, Case No. CV09 05718, *NETLIST, INC. vs. GOOGLE, INC.*, Plaintiff Netlist, Inc.'s Reply Claim Construction Brief, filed Aug. 16, 2010 in 17 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Complaint for Declaratory Relief, filed Aug. 29, 2008 in 49 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Netlist, Inc.'s Answer To Complaint And Counterclaims, filed Nov. 18, 2008 in 9 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Plaintiff Google's Reply To Counterclaims, filed Dec. 8, 2008 in 4 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Joint Claim Construction And Prehearing Statement, filed Jun. 12, 2009 in 5 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Exhibit A To Joint Claim Construction And Prehearing Statement, filed Jun. 12, 2009 in 2 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Exhibit B To Joint Claim Construction And Prehearing Statement, filed Jun. 12, 2009 in 36 pages.

US 7,881,150 B2

Page 5

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Attachment 1 To Exhibit B To Joint Claim Construction And Prehearing Statement, filed Jun. 12, 2009 in 7 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Attachment 2 To Exhibit B To Joint Claim Construction And Prehearing Statement, filed Jun. 12, 2009 in 12 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Defendant Netlist, Inc.'s Opening Claim Construction Brief, filed Jul. 29, 2009 in 21 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, [Redacted] Google Inc.'s Responsive Claim Construction Brief, filed Aug. 25, 2009 in 30 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Appendix 1 To Google's Responsive Claim Construction Brief, filed Aug. 25, 2009 in 4 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Defendant Netlist, Inc.'s Claim Construction Reply Brief, filed Sep. 22, 2009 in 19 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Stipulation Re: Additional Agreed-Upon Claim Constructions, filed Oct. 28, 2009 in 3 pages.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Amended Exhibit A To Joint Claim Construction And Prehearing Statement, filed Oct. 28, 2009 in 1 page.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Order Re Claim Construction, filed Nov. 16, 2009 in 1 page.

U.S. District Court Northern District of California, Case No. CV08 04144, *GOOGLE INC. vs. NETLIST, INC.*, Defendant Netlist, Inc.'s Opposition To Google Inc's Motion For Summary Judgement Of Invalidity, filed Jul. 6, 2010 in 13 pages.

Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,577 for U.S. Pat. No. 7,289,386 filed Oct. 20, 2010.

Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,578 for U.S. Pat. No. 7,619,912 filed Oct. 20, 2010.

Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,579 for U.S. Pat. No. 7,619,912 filed Oct. 21, 2010.

"DDR SDRAM RDIMM Features," Micron Technology, Inc., 2002. Freedman, Alan. "The Computer Glossary," The Complete Illustrated Dictionary, American Management Association, 2001.

Jedec Standard No. 21-C, "PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification," Revision 1.3, Jan. 2002.

Jedec Standard, "Definition of the SSTV16859 2.5 V 13-Bit to 26-Bit SSTL_2 Registered Buffer for Stacked DDR DIMM Applications," JESD82-4B, May 2003.

"Quad Band Memory (QBM™): DDR 200/266/333 devices producing DDR 400/533/667," Platform Conference, Jan. 23-24, 2002.

Portion of Request for *Inter Partes* Reexamination of U.S. Patent No. 7,289,386, corresponding to Reexam U.S. Appl. No. 95/000,577, in 184 pages.

Portion of Request for *Inter Partes* Reexamination of U.S. Patent No. 7,619,912, corresponding to Reexam U.S. Appl. No. 95/000,578, in 66 pages.

Portion of Request for *Inter Partes* Reexamination of U.S. Patent No. 7,619,912, corresponding to Reexam U.S. Appl. No. 95/000,579, in 32 pages.

US 6,438,062, 08/2002, Curtis et al. (withdrawn)

* cited by examiner

U.S. Patent

Feb. 1, 2011

Sheet 1 of 23

US 7,881,150 B2

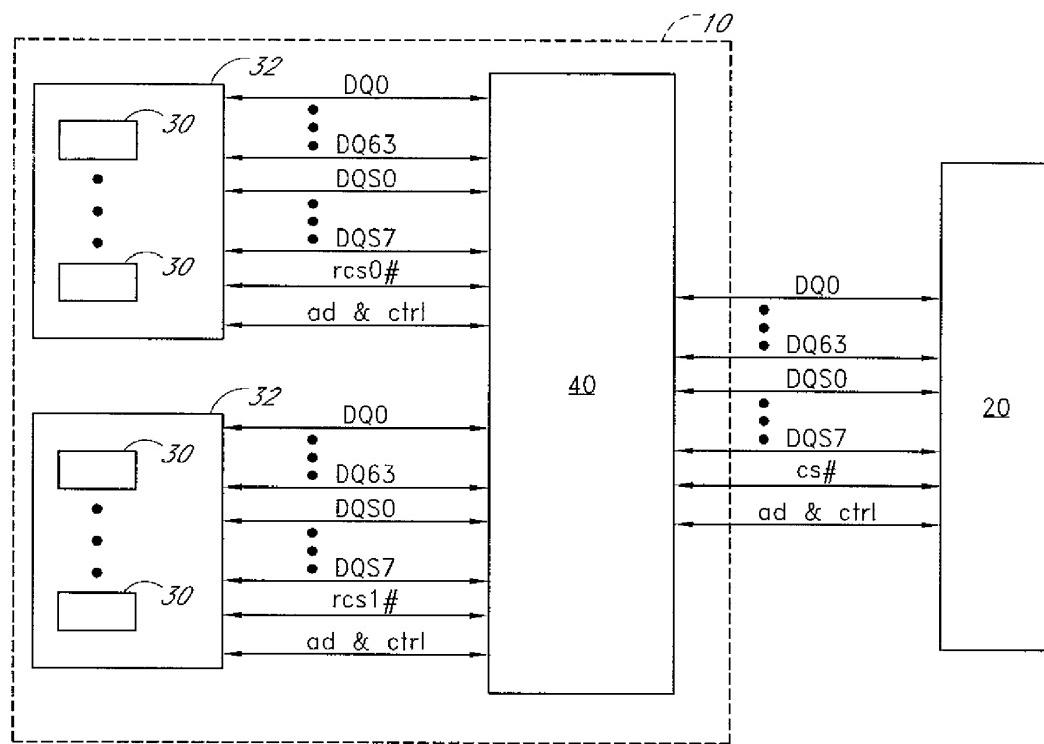


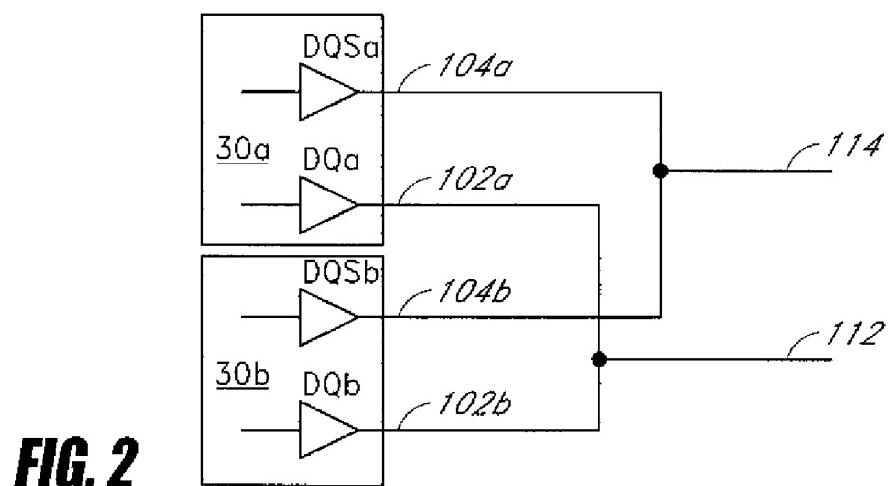
FIG. 1

U.S. Patent

Feb. 1, 2011

Sheet 2 of 23

US 7,881,150 B2



U.S. Patent

Feb. 1, 2011

Sheet 3 of 23

US 7,881,150 B2

FIG. 3A

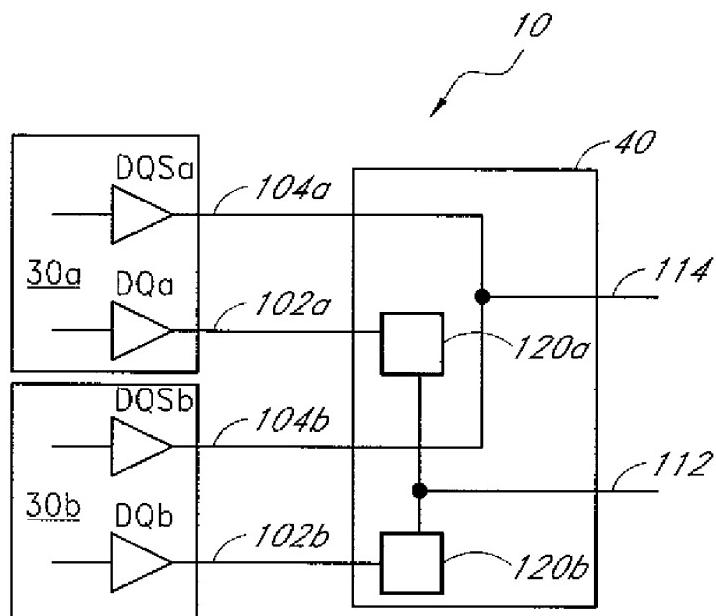
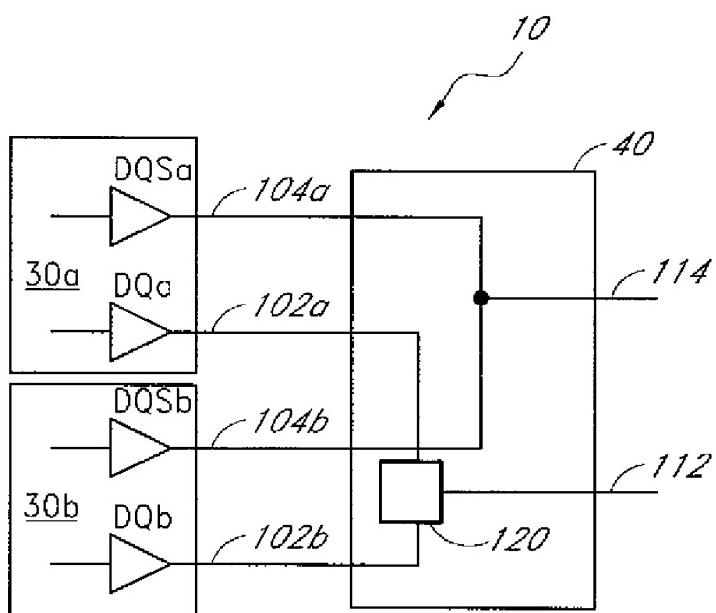


FIG. 3B

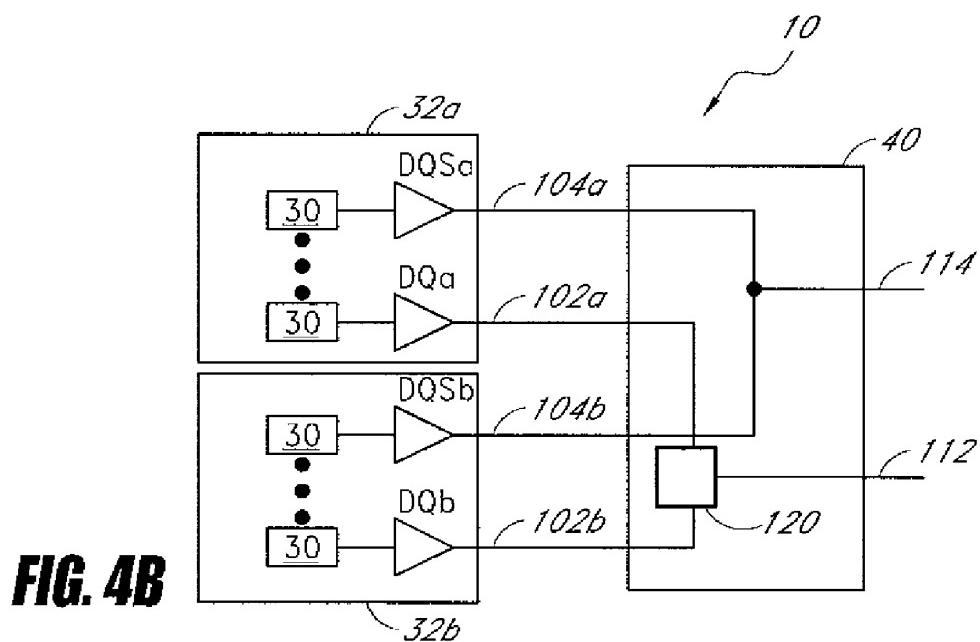
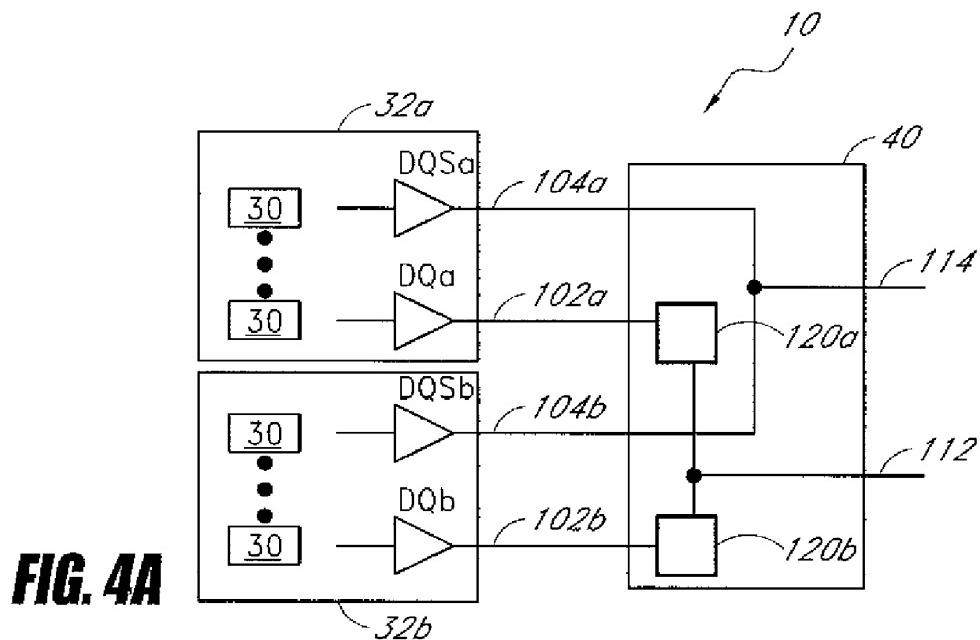


U.S. Patent

Feb. 1, 2011

Sheet 4 of 23

US 7,881,150 B2

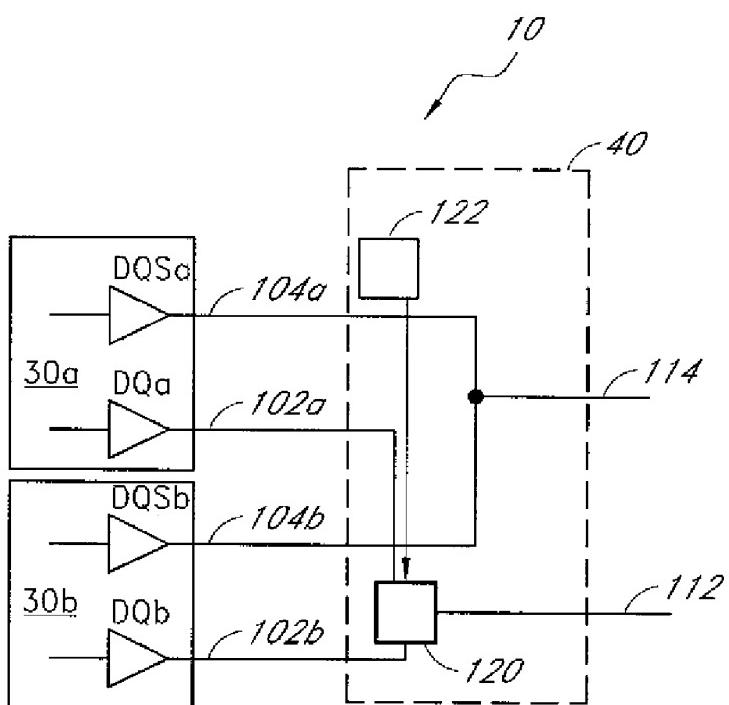
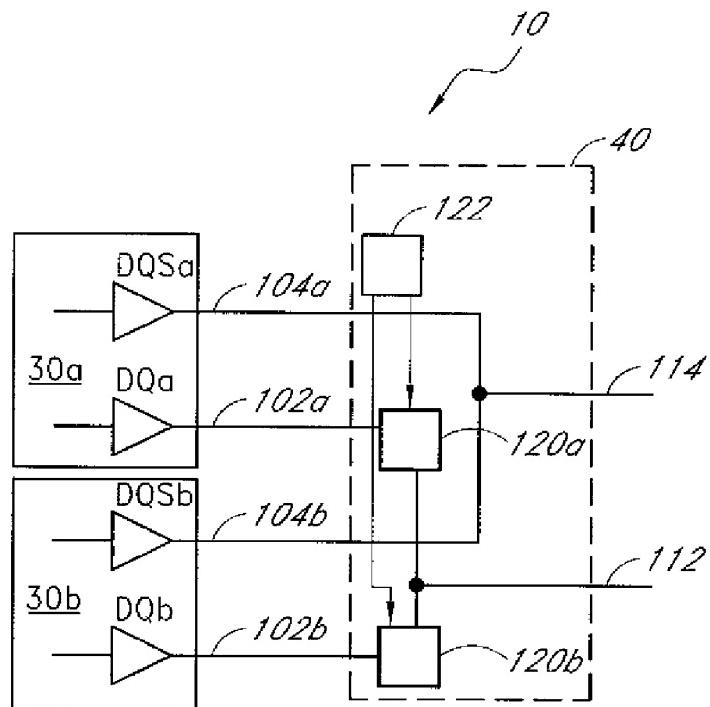


U.S. Patent

Feb. 1, 2011

Sheet 5 of 23

US 7,881,150 B2



U.S. Patent

Feb. 1, 2011

Sheet 6 of 23

US 7,881,150 B2

FIG. 5C

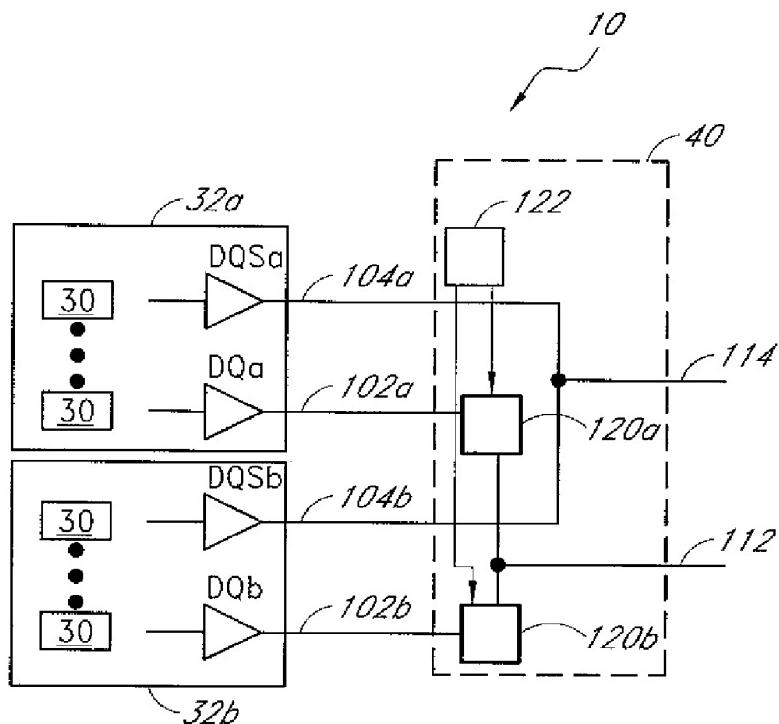
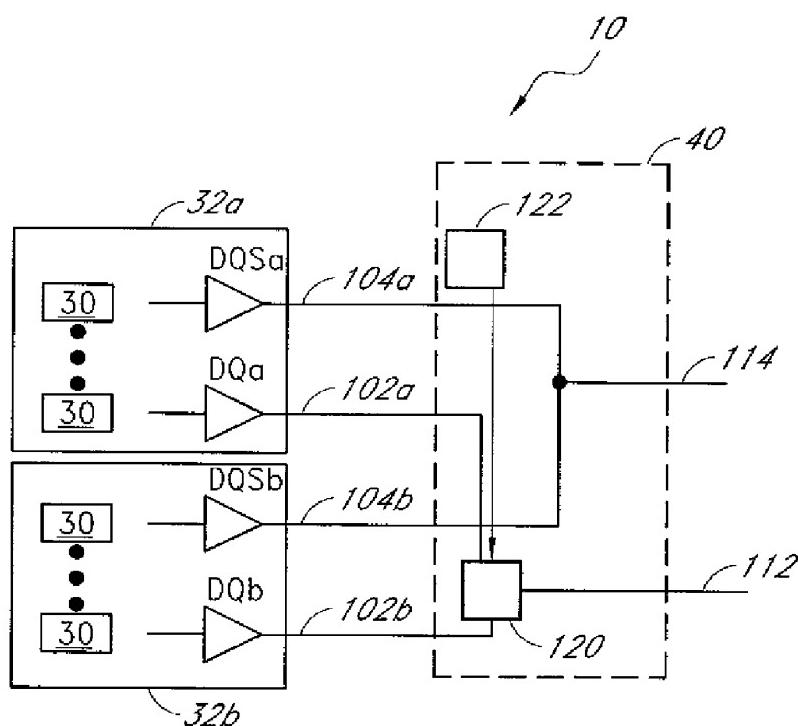


FIG. 5D



U.S. Patent

Feb. 1, 2011

Sheet 7 of 23

US 7,881,150 B2

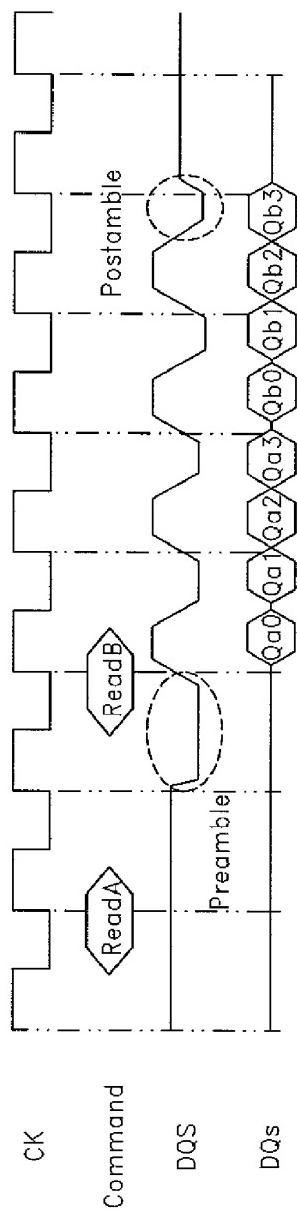


FIG. 6A

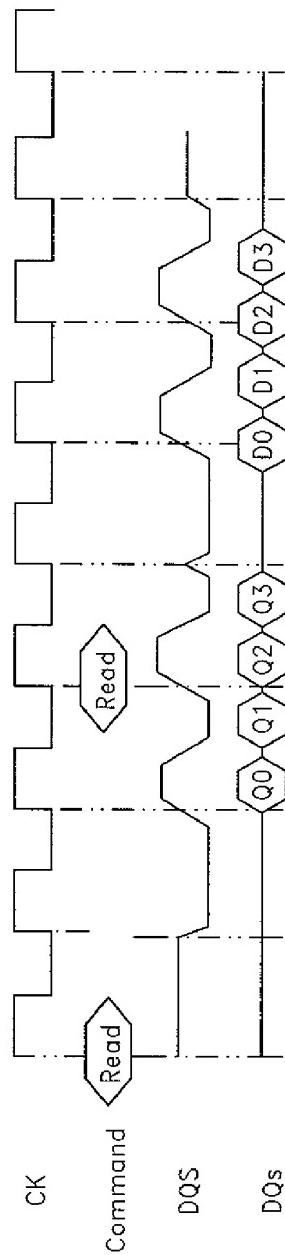


FIG. 6B

U.S. Patent

Feb. 1, 2011

Sheet 8 of 23

US 7,881,150 B2

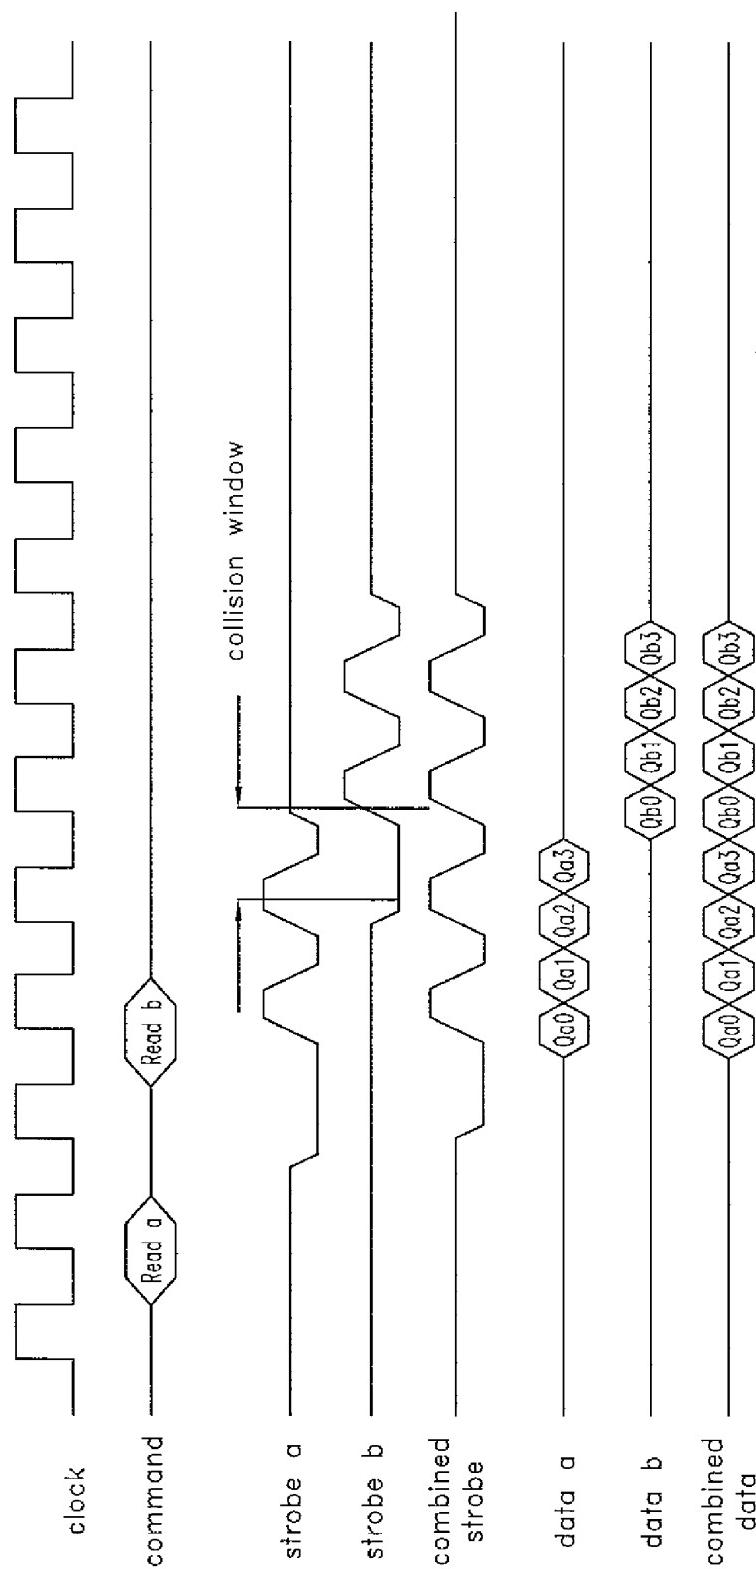


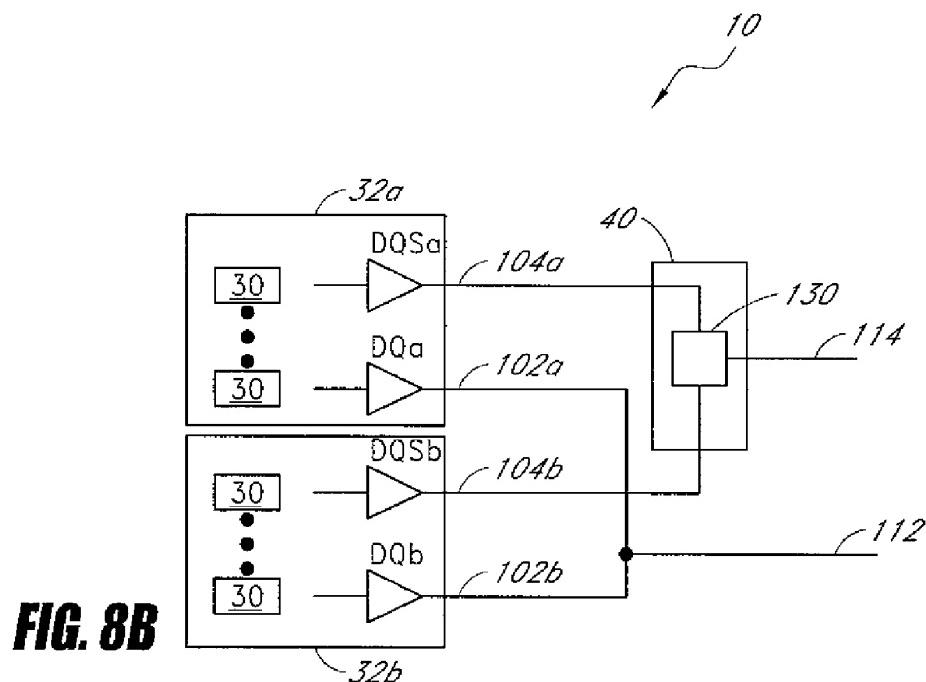
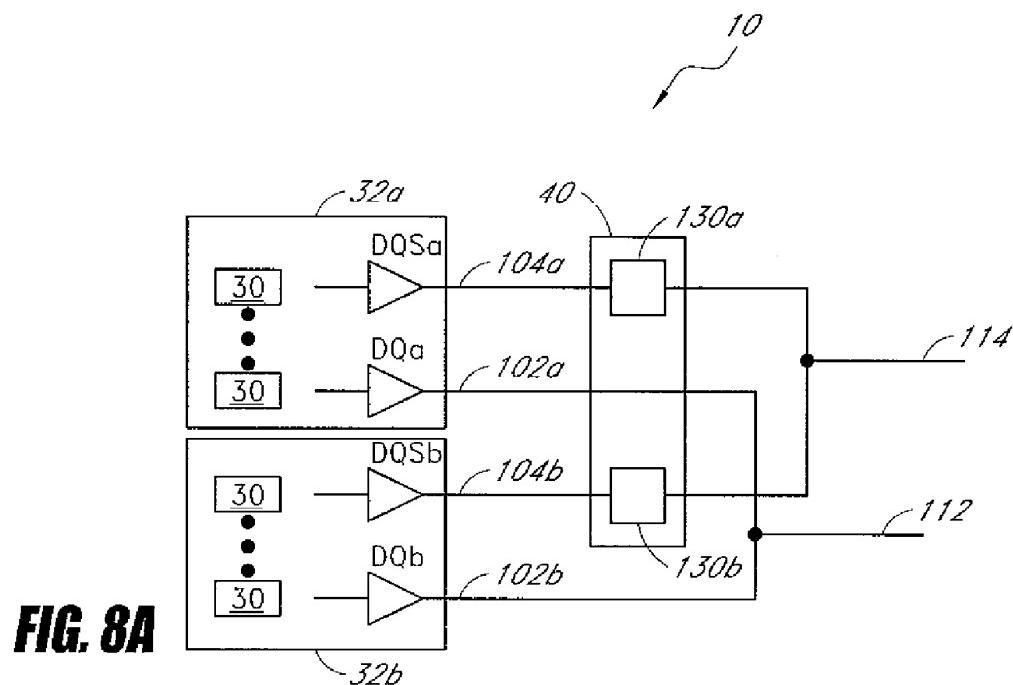
FIG. 7

U.S. Patent

Feb. 1, 2011

Sheet 9 of 23

US 7,881,150 B2



U.S. Patent

Feb. 1, 2011

Sheet 10 of 23

US 7,881,150 B2

FIG. 8C

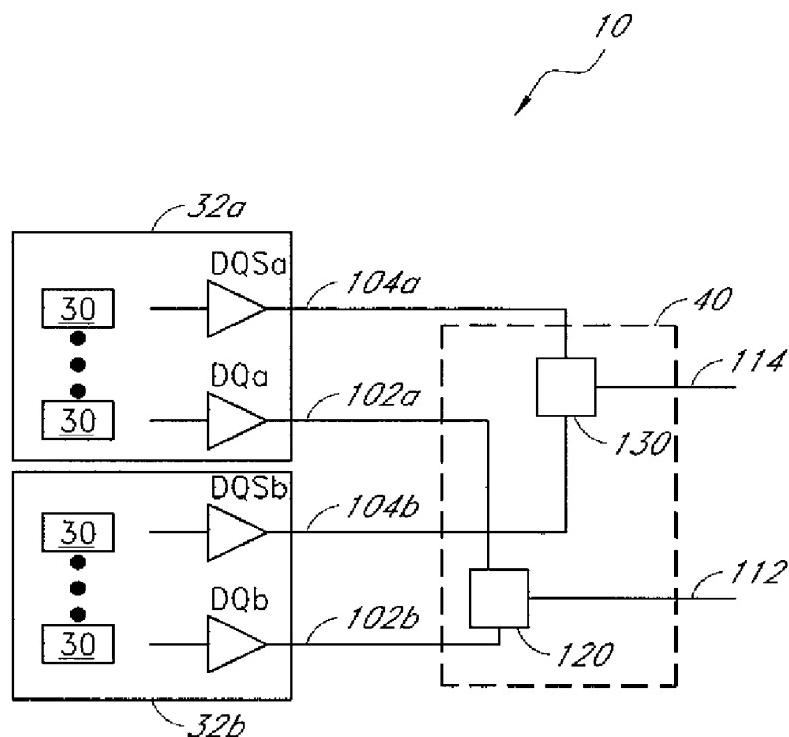
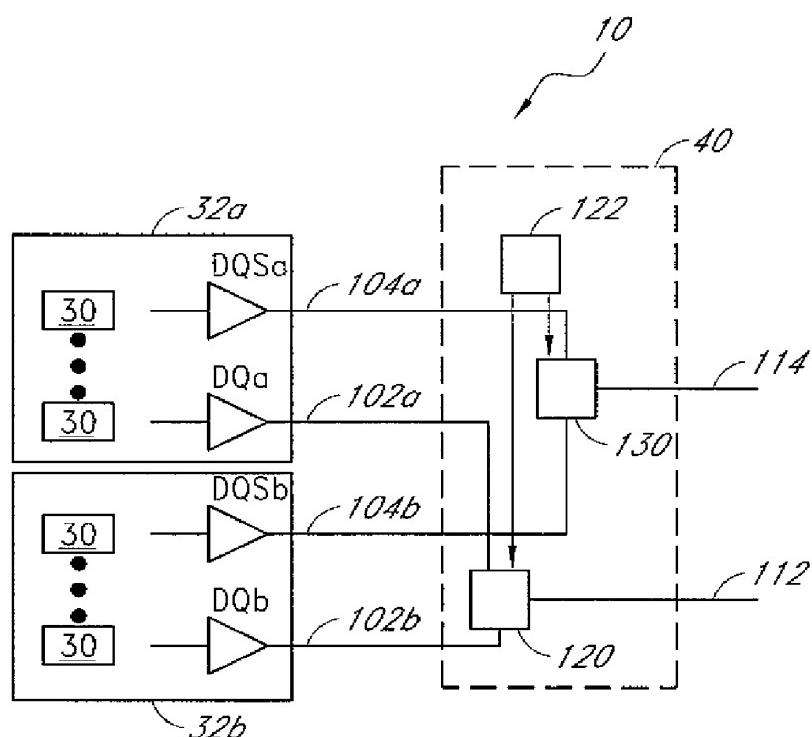


FIG. 8D



U.S. Patent

Feb. 1, 2011

Sheet 11 of 23

US 7,881,150 B2

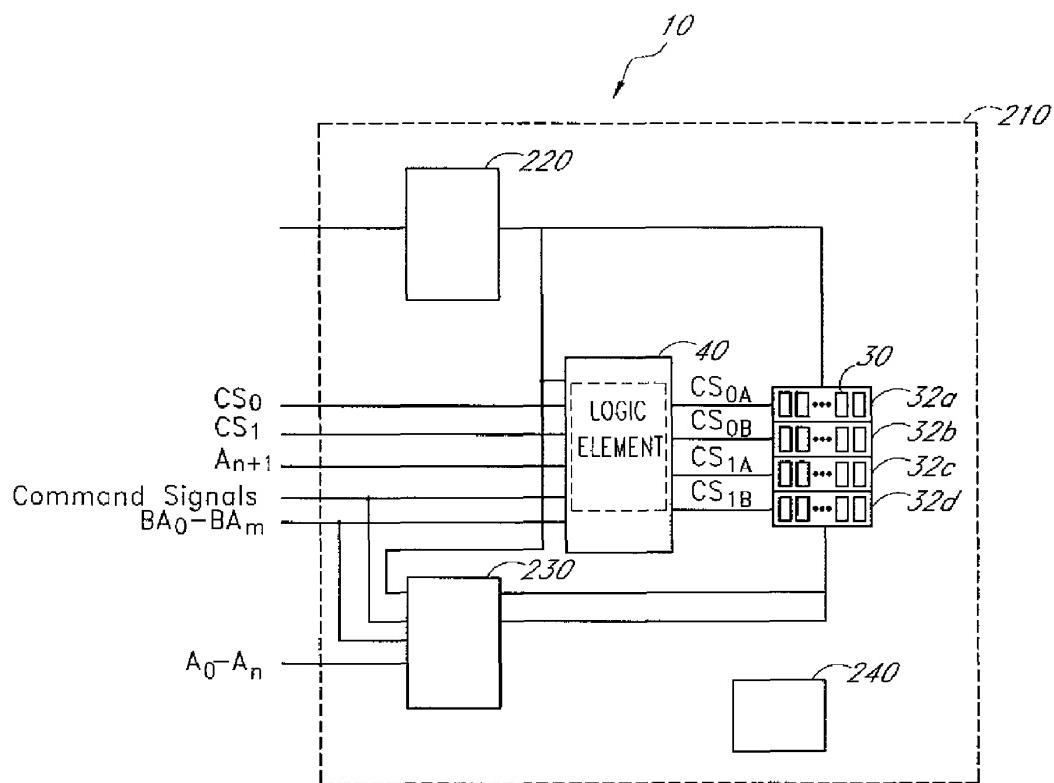


FIG. 9A

U.S. Patent

Feb. 1, 2011

Sheet 12 of 23

US 7,881,150 B2

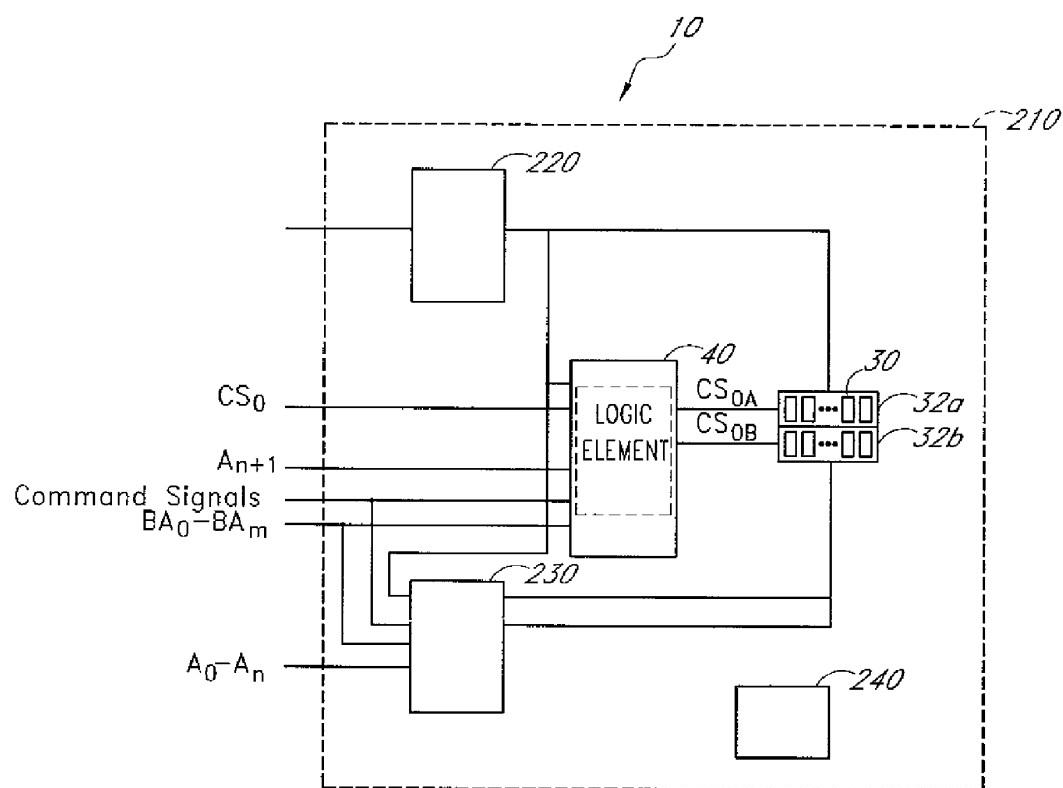


FIG. 9B

U.S. Patent

Feb. 1, 2011

Sheet 13 of 23

US 7,881,150 B2

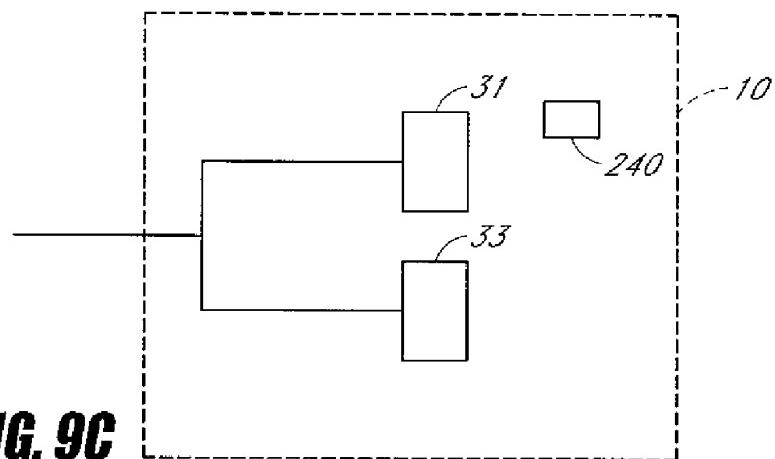


FIG. 9C

U.S. Patent

Feb. 1, 2011

Sheet 14 of 23

US 7,881,150 B2

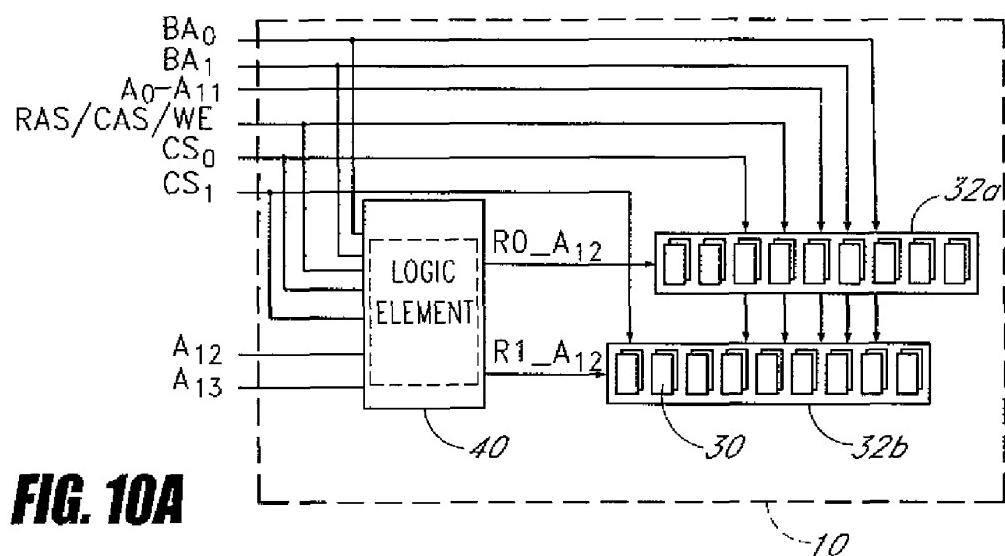


FIG. 10A

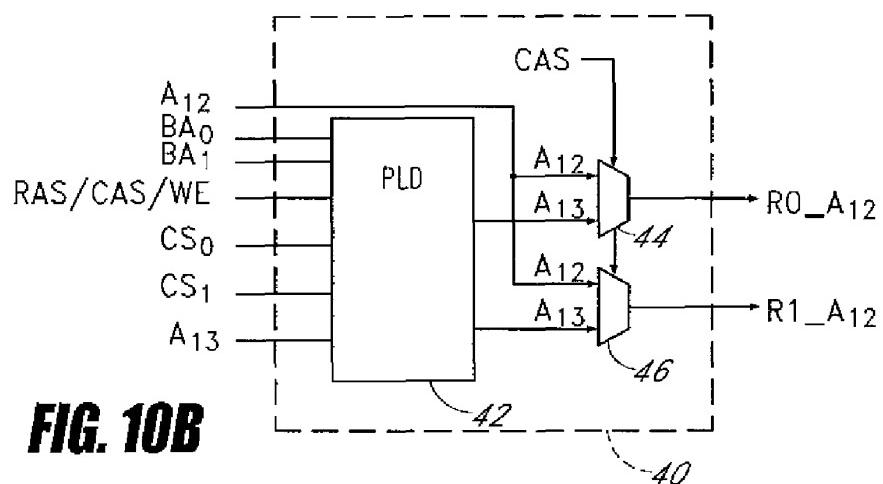


FIG. 10B

U.S. Patent

Feb. 1, 2011

Sheet 15 of 23

US 7,881,150 B2

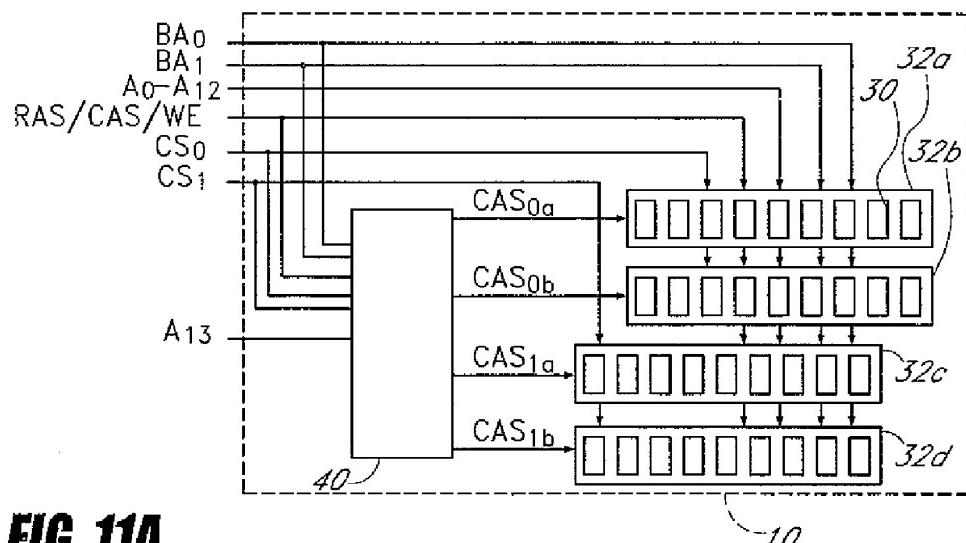


FIG. 11A

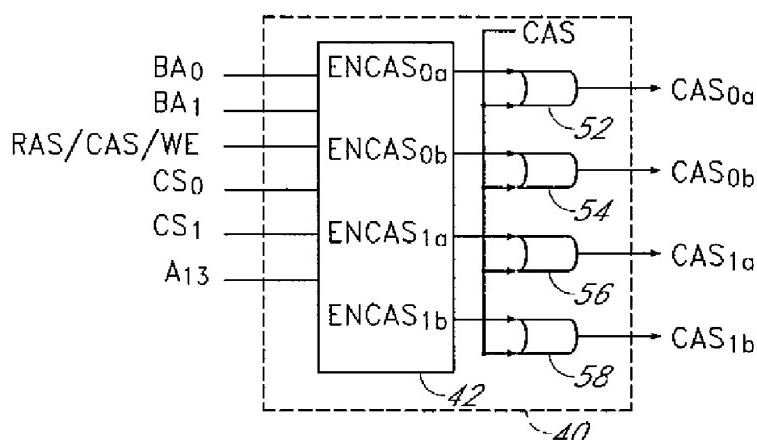


FIG. 11B

U.S. Patent

Feb. 1, 2011

Sheet 16 of 23

US 7,881,150 B2

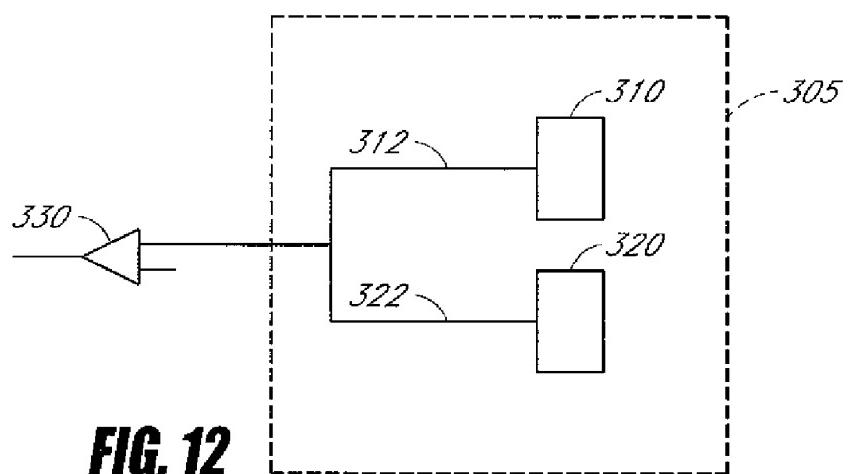


FIG. 12

U.S. Patent

Feb. 1, 2011

Sheet 17 of 23

US 7,881,150 B2

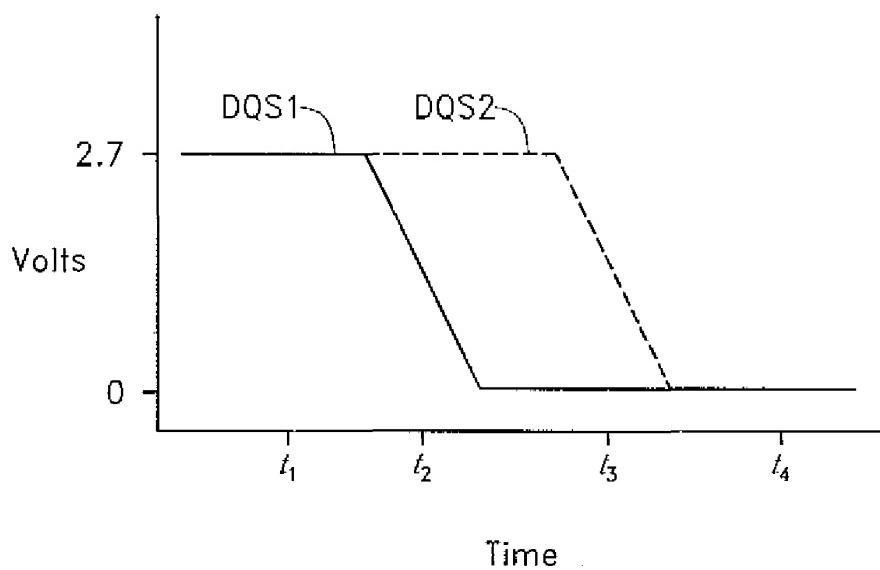


FIG. 13

U.S. Patent

Feb. 1, 2011

Sheet 18 of 23

US 7,881,150 B2

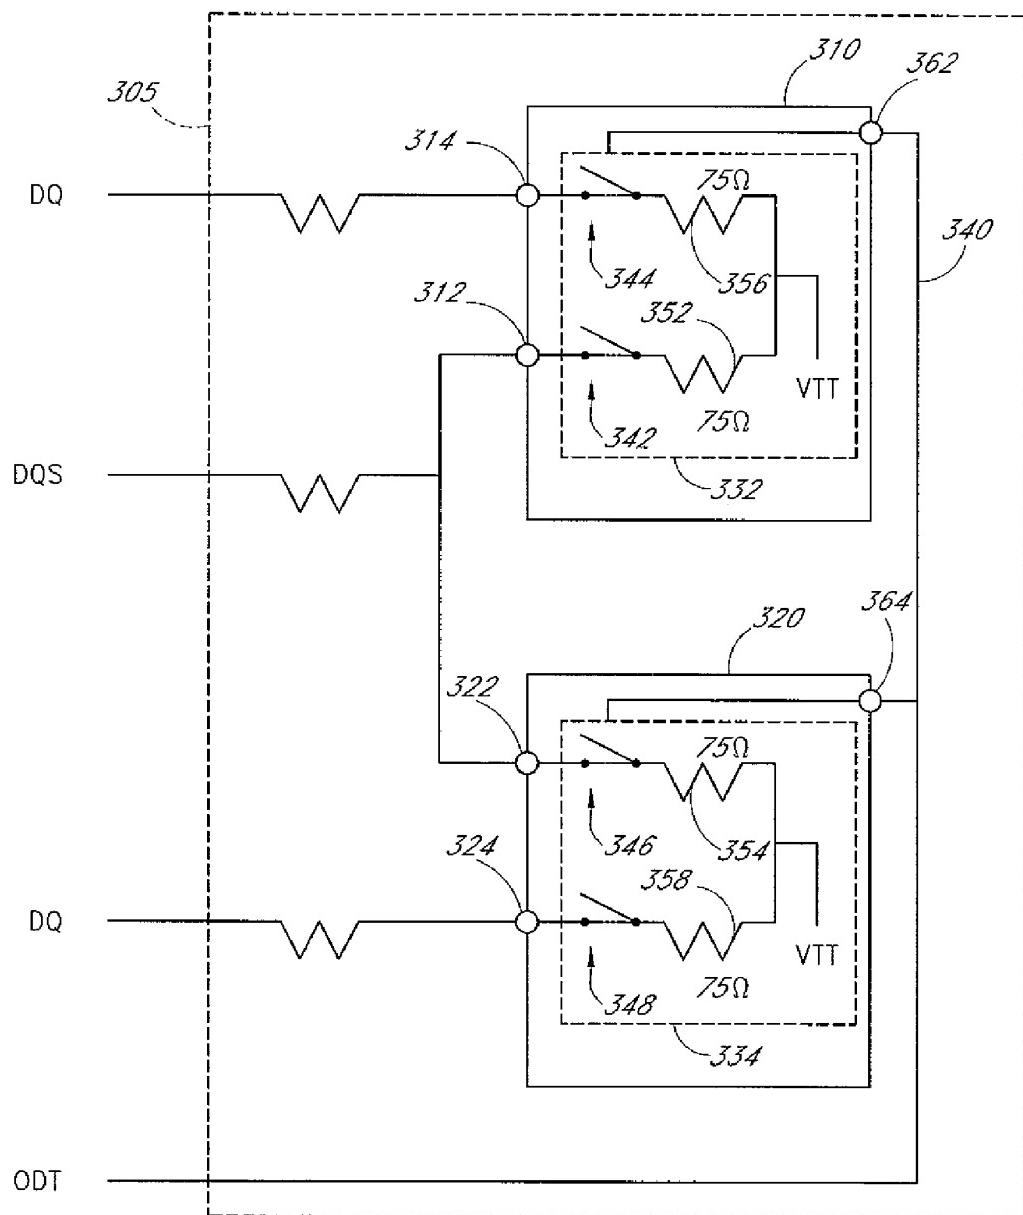


FIG. 14

U.S. Patent

Feb. 1, 2011

Sheet 19 of 23

US 7,881,150 B2

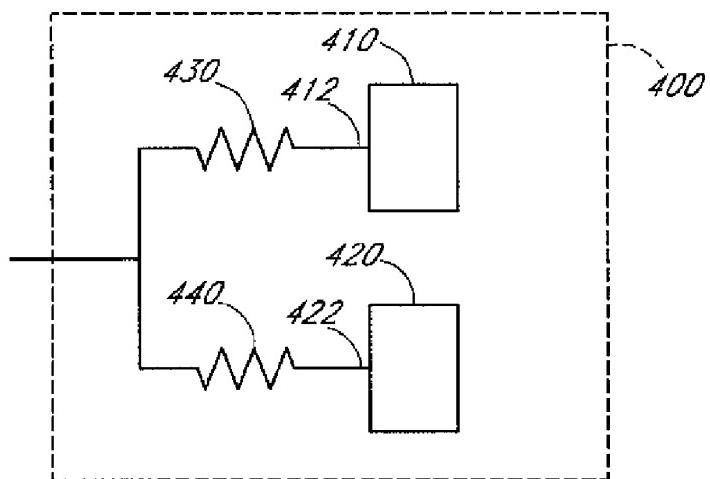


FIG. 15

U.S. Patent

Feb. 1, 2011

Sheet 20 of 23

US 7,881,150 B2

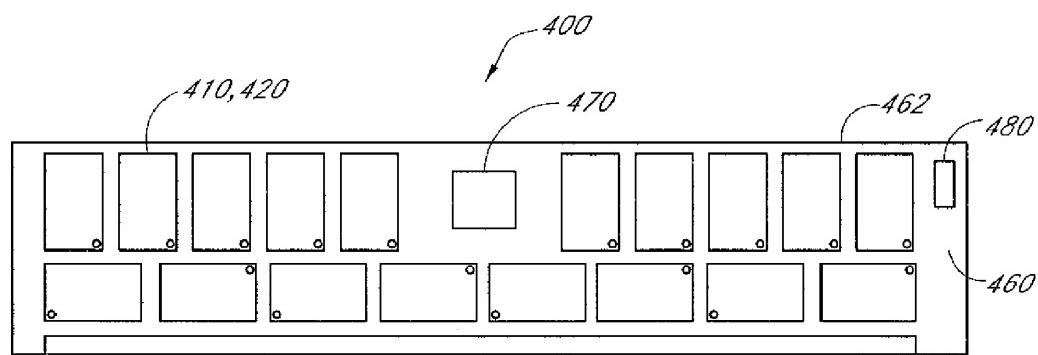


FIG. 16A

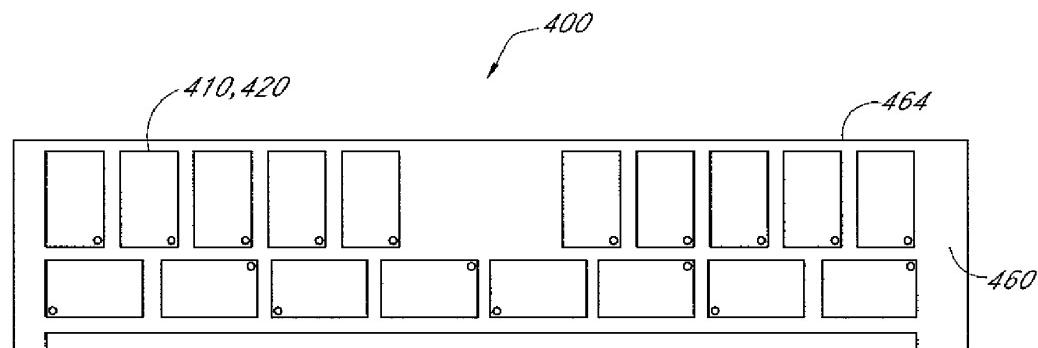


FIG. 16B

U.S. Patent

Feb. 1, 2011

Sheet 21 of 23

US 7,881,150 B2

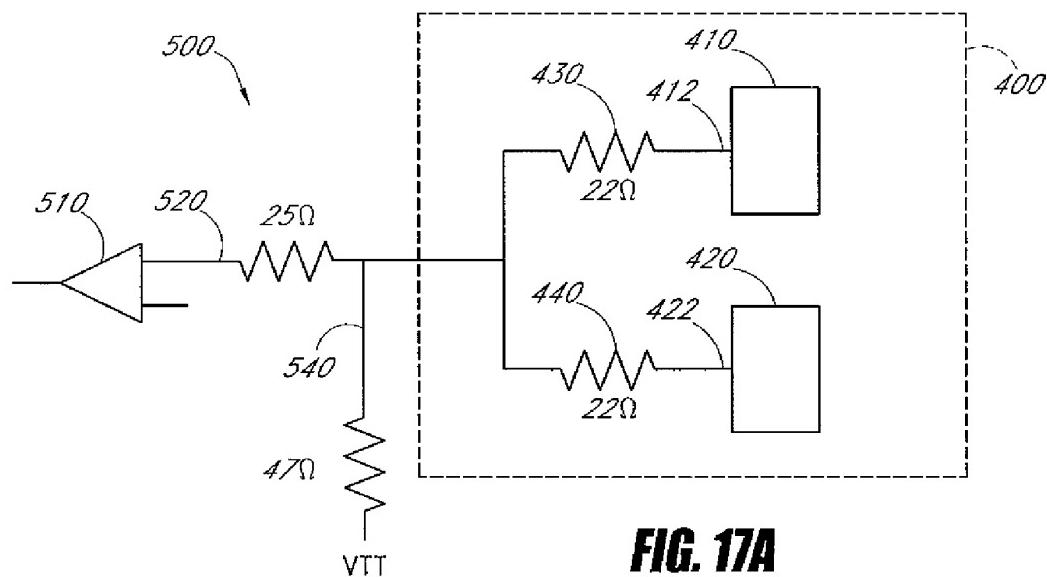


FIG. 17A

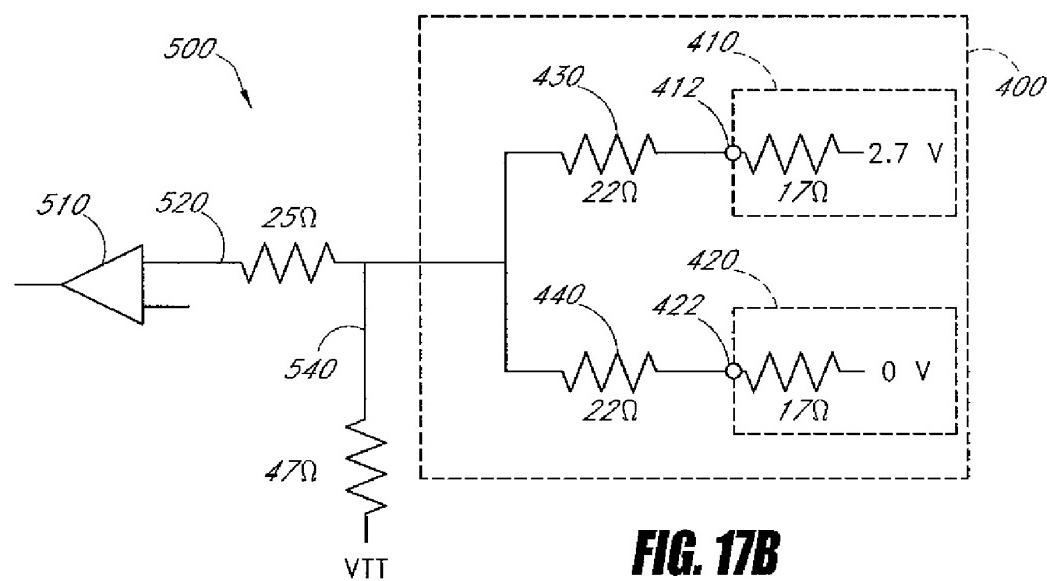


FIG. 17B

U.S. Patent

Feb. 1, 2011

Sheet 22 of 23

US 7,881,150 B2

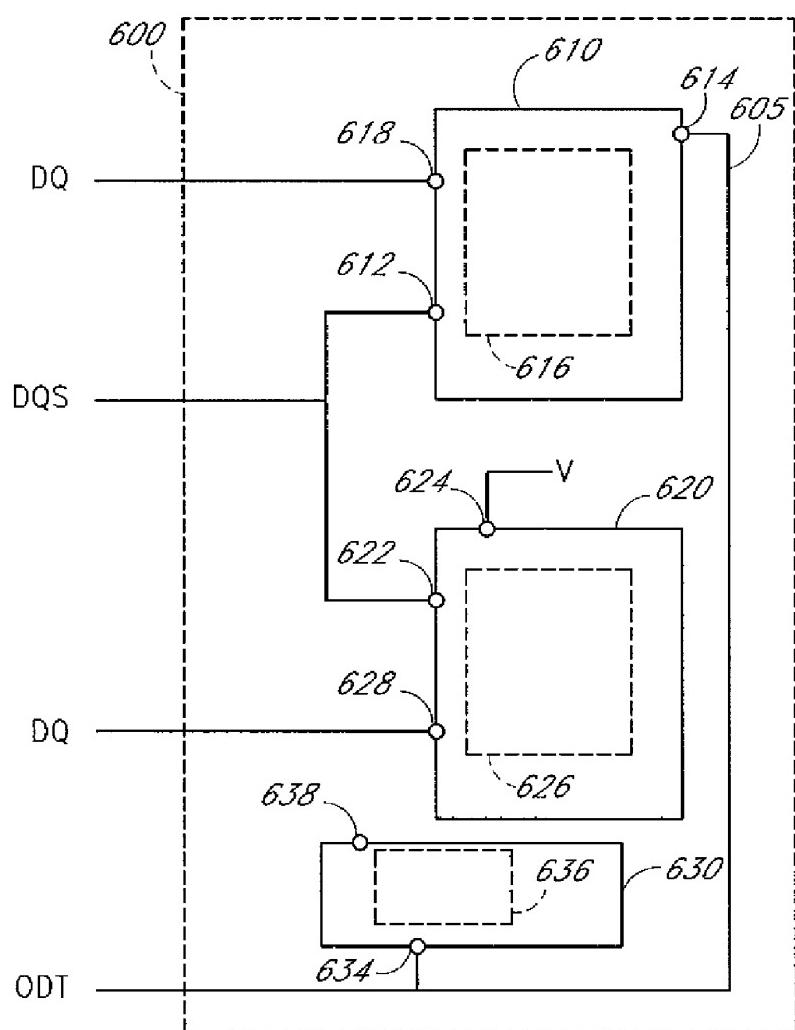


FIG. 18

U.S. Patent

Feb. 1, 2011

Sheet 23 of 23

US 7,881,150 B2

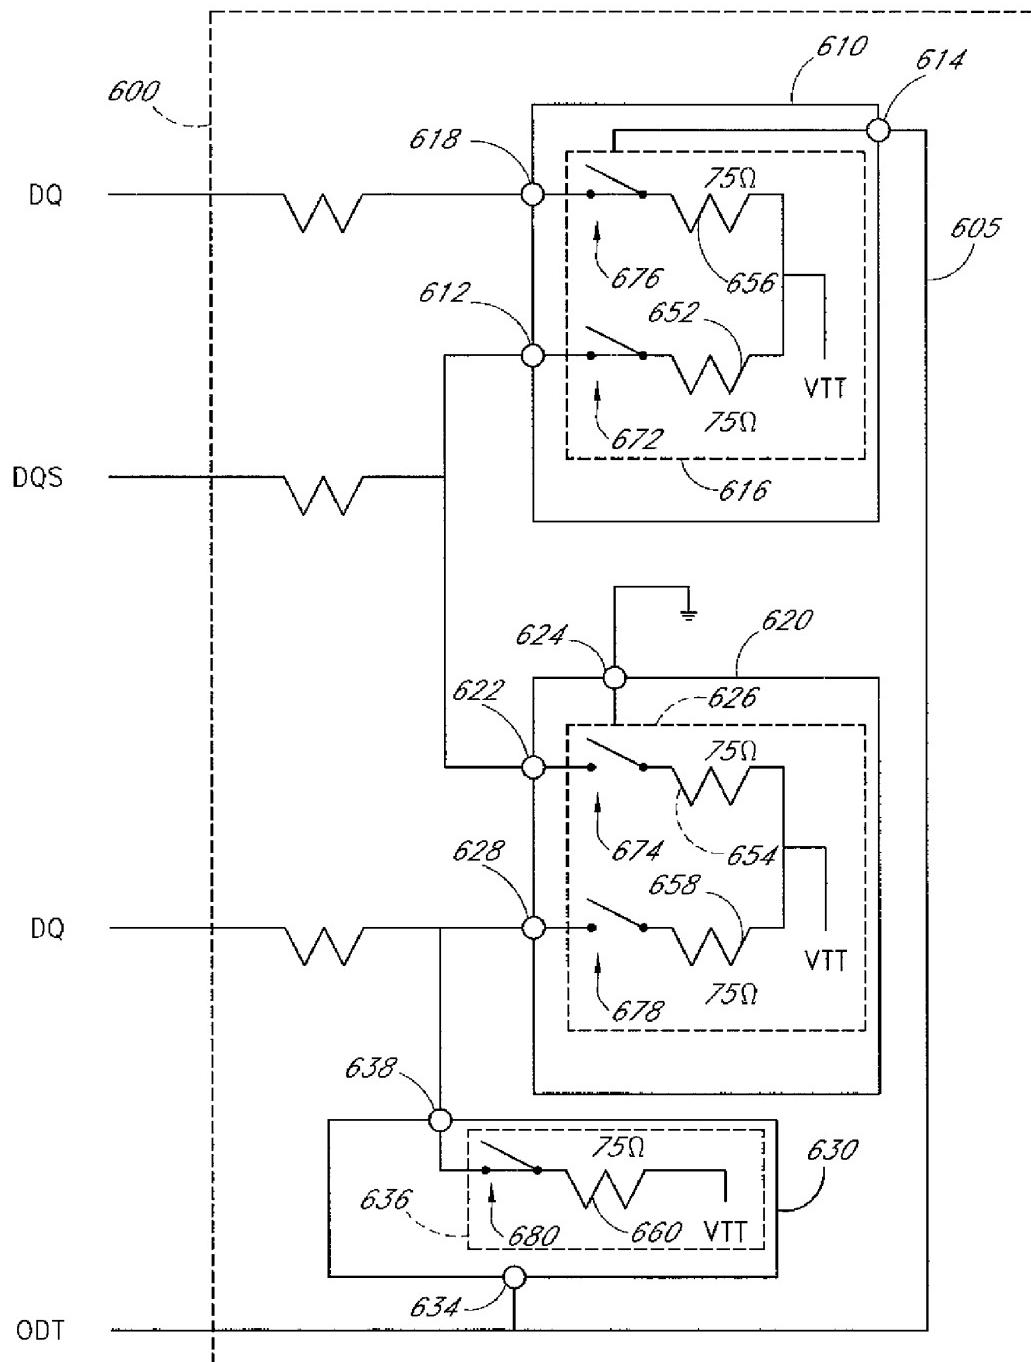


FIG. 19

US 7,881,150 B2

1**CIRCUIT PROVIDING LOAD ISOLATION
AND MEMORY DOMAIN TRANSLATION
FOR MEMORY MODULE****CROSS-REFERENCE TO RELATED
APPLICATIONS**

The present application is continuation of U.S. patent application Ser. No. 12/408,652, filed Mar. 20, 2009, which is a continuation of U.S. patent application Ser. No. 11/335,875, filed Jan. 19, 2006, which claims the benefit of U.S. Provisional Appl. No. 60/645,087, filed Jan. 19, 2005 and which is a continuation-in-part of U.S. patent application Ser. No. 11/173,175, filed Jul. 1, 2005, which claims the benefit of U.S. Provisional Appl. No. 60/588,244, filed Jul. 15, 2004 and which is a continuation-in-part of U.S. patent application Ser. No. 11/075,395, filed Mar. 7, 2005, which claims the benefit of U.S. Provisional Appl. No. 60/550,668, filed Mar. 5, 2004, U.S. Provisional Appl. No. 60/575,595, filed May 28, 2004, and U.S. Provisional Appl. No. 60/590,038, filed Jul. 21, 2004. U.S. patent application Ser. Nos. 12/408,652, 11/335,875, 11/173,175, and 11/075,395 and U.S. Provisional Appl. Nos. 60/550,668, 60/575,595, 60/590,038, 60/588,244, and 60/645,087 are each incorporated in their entirety by reference herein.

BACKGROUND OF THE INVENTION**1. Field of the Invention**

The present invention relates generally to memory modules of a computer system, and more specifically to devices and methods for improving the performance, the memory capacity, or both, of memory modules.

2. Description of the Related Art

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the processor of the computer system. Memory modules typically have a memory configuration with a unique combination of rows, columns, and banks which result in a total memory capacity for the memory module.

For example, a 512-Megabyte memory module (termed a “512-MB” memory module, which actually has 2^{29} or 536,870,912 bytes of capacity) will typically utilize eight 512-Megabit DRAM devices (each identified as a “512-Mb” DRAM device, each actually having 2^{29} or 536,870,912 bits of capacity). The memory cells (or memory locations) of each 512-Mb DRAM device can be arranged in four banks, with each bank having an array of 2^{24} (or 16,777,216) memory locations arranged as 2^{13} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 64M 8-bit-wide memory locations (actually with four banks of 2^{27} or 134,217,728 one-bit memory cells arranged to provide a total of 2^{26} or 67,108,864 memory locations with 8 bits each) are identified as having a “64 Mb×8” or “64M×8-bit” configuration, or as having a depth of 64M and a bit width of 8. Furthermore, certain commercially-available 512-MB memory modules are termed to have a “64M×8-byte” configuration or a “64M×64-bit” configuration with a depth of 64M and a width of 8 bytes or 64 bits.

Similarly, a 1-Gigabyte memory module (termed a “1-GB” memory module, which actually has 2^{30} or 1,073,741,824 bytes of capacity) can utilize eight 1-Gigabit DRAM devices (each identified as a “1-Gb” DRAM device, each actually having 2^{30} or 1,073,741,824 bits of capacity). The memory

2

locations of each 1-Gb DRAM device can be arranged in four banks, with each bank having an array of memory locations with 2^{14} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 128M 8-bit-wide memory locations (actually with a total of 2^{27} or 134,217,728 memory locations with 8 bits each) are identified as having a “128 Mb×8” or “128M×8-bit” configuration, or as having a depth of 128M and a bit width of 8. Furthermore, certain commercially-available 1-GB memory modules are identified as having a “128M×8-byte” configuration or a “128M×64-bit” configuration with a depth of 128M and a width of 8 bytes or 64 bits.

The commercially-available 512-MB (64M×8-byte) memory modules and the 1-GB (128M×8-byte) memory modules described above are typically used in computer systems (e.g., personal computers) which perform graphics applications since such “×8” configurations are compatible with data mask capabilities often used in such graphics applications. Conversely, memory modules with “×4” configurations are typically used in computer systems such as servers which are not as graphics-intensive. Examples of such commercially available “×4” memory modules include, but are not limited to, 512-MB (128M×4-byte) memory modules comprising eight 512-Mb (128 Mb×4) memory devices.

25 The DRAM devices of a memory module are generally arranged as ranks or rows of memory, each rank of memory generally having a bit width. For example, a memory module in which each rank of the memory module is 64 bits wide is described as having an “×64” organization. Similarly, a memory module having 72-bit-wide ranks is described as having an “×72” organization.

30 The memory capacity of a memory module increases with the number of memory devices. The number of memory devices of a memory module can be increased by increasing the number of memory devices per rank or by increasing the number of ranks. For example, a memory module with four ranks has double the memory capacity of a memory module with two ranks and four times the memory capacity of a memory module with one rank. Rather than referring to the 35 memory capacity of the memory module, in certain circumstances, the memory density of the memory module is referred to instead.

35 During operation, the ranks of a memory module are selected or activated by address and command signals that are received from the processor. Examples of such address and command signals include, but are not limited to, rank-select signals, also called chip-select signals. Most computer and server systems support one-rank and two-rank memory modules. By only supporting one-rank and two-rank memory 40 modules, the memory density that can be incorporated in each memory slot is limited.

Various aspects of the design of a memory module impose 45 limitations on the size of the memory arrays of the memory module. Certain such aspects are particularly important for memory modules designed to operate at higher frequencies. Examples of such aspects include, but are not limited to, memory device (e.g., chip) densities, load fan-out, signal integrity, available rank selects, power dissipation, and thermal profiles.

SUMMARY OF THE INVENTION

In certain embodiments, a memory module comprises a plurality of memory devices and a circuit. Each memory device has a corresponding load. The circuit is electrically coupled to the plurality of memory devices and is configured to be electrically coupled to a memory controller of a com-

US 7,881,150 B2

3

puter system. The circuit selectively isolates one or more of the loads of the memory devices from the computer system. The circuit comprises logic which translates between a system memory domain of the computer system and a physical memory domain of the memory module.

In certain embodiments, a method of using a memory module with a computer system is provided. The method comprises providing a memory module having a plurality of memory devices and a circuit. The circuit is electrically coupled to the plurality of memory devices. The circuit is configured to be electrically coupled to a computer system. Each memory device has a corresponding load. The method further comprises electrically coupling the memory module to the computer system. The method further comprises activating the circuit to selectively isolate at least one of the loads of the memory devices from the computer system and to translate between a system memory domain of the computer system and a physical memory domain of the memory module.

In certain embodiments, a memory module is connectable to a computer system. The memory module comprises a first memory device having a first data signal line and a first data strobe signal line. The memory module further comprises a second memory device having a second data signal line and a second data strobe signal line. The memory module further comprises a common data signal line connectable to the computer system. The memory module further comprises a device electrically coupled to the first data signal line, to the second data signal line, and to the common data signal line. The device selectively electrically couples the first data signal line to the common data signal line and selectively electrically couples the second data signal line to the common data signal line. The device comprises logic which translates between a system memory domain of the computer system and a physical memory domain of the memory module.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 schematically illustrates an example memory module in accordance with certain embodiments described herein.

FIG. 2 schematically illustrates a circuit diagram of two memory devices of a conventional memory module.

FIGS. 3A and 3B schematically illustrate example memory modules having a circuit which selectively isolates one or both of the DQ data signal lines of the two memory devices from the computer system in accordance with certain embodiments described herein.

FIGS. 4A and 4B schematically illustrate example memory modules having a circuit which selectively isolates one or both of the DQ data signal lines of the two ranks of memory devices from the computer system in accordance with certain embodiments described herein.

FIGS. 5A-5D schematically illustrate example memory modules having a circuit comprising a logic element and one or more switches operatively coupled to the logic element in accordance with certain embodiments described herein.

FIG. 6A shows an exemplary timing diagram of a gapless read burst for a back-to-back adjacent read condition from one memory device.

FIG. 6B shows an exemplary timing diagram with an extra clock cycle between successive read commands issued to different memory devices for successive read accesses from different memory devices.

FIG. 7 shows an exemplary timing diagram in which the last data strobe of memory device "a" collides with the preamble time interval of the data strobe of memory device "b."

4

FIGS. 8A-8D schematically illustrate circuit diagrams of example memory modules comprising a circuit which multiplexes the DQS data strobe signal lines from one another in accordance with certain embodiments described herein.

FIG. 9A schematically illustrates an example memory module with four ranks of memory devices compatible with certain embodiments described herein.

FIG. 9B schematically illustrates an example memory module with two ranks of memory devices compatible with certain embodiments described herein.

FIG. 9C schematically illustrates another example memory module in accordance with certain embodiments described herein.

FIG. 10A schematically illustrates an exemplary memory module which doubles the rank density in accordance with certain embodiments described herein.

FIG. 10B schematically illustrates an exemplary circuit compatible with embodiments described herein.

FIG. 11A schematically illustrates an exemplary memory module which doubles number of ranks in accordance with certain embodiments described herein.

FIG. 11B schematically illustrates an exemplary circuit compatible with embodiments described herein.

FIG. 12 schematically illustrates an exemplary memory module in which a data strobe (DQS) pin of a first memory device is electrically connected to a DQS pin of a second memory device while both DQS pins are active.

FIG. 13 is an exemplary timing diagram of the voltages applied to the two DQS pins due to non-simultaneous switching.

FIG. 14 schematically illustrates another exemplary memory module in which a DQS pin of a first memory device is connected to a DQS pin of a second memory device.

FIG. 15 schematically illustrates an exemplary memory module in accordance with certain embodiments described herein.

FIGS. 16A and 16B schematically illustrate a first side and a second side, respectively, of a memory module with eighteen 64Mx4 bit, DDR-1 SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board.

FIGS. 17A and 17B schematically illustrate an exemplary embodiment of a memory module in which a first resistor and a second resistor are used to reduce the current flow between the first DQS pin and the second DQS pin.

FIG. 18 schematically illustrates another exemplary memory module compatible with certain embodiments described herein.

FIG. 19 schematically illustrates a particular embodiment of the memory module schematically illustrated by FIG. 18.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

55 Load Isolation

FIG. 1 schematically illustrates an example memory module 10 compatible with certain embodiments described herein. The memory module 10 is connectable to a memory controller 20 of a computer system (not shown). The memory module 10 comprises a plurality of memory devices 30, each memory device 30 having a corresponding load. The memory module 10 further comprises a circuit 40 electrically coupled to the plurality of memory devices 30 and configured to be electrically coupled to the memory controller 20 of the computer system. The circuit 40 selectively isolates one or more of the loads of the memory devices from the computer system. The circuit 40 comprises logic which translates between a system memory domain of the computer system and a physical memory domain of the memory module 10.

US 7,881,150 B2

5

system memory domain of the computer system and a physical memory domain of the memory module 10.

As used herein, the term “load” is a broad term which includes, without limitation, electrical load, such as capacitive load, inductive load, or impedance load. As used herein, the term “isolation” is a broad term which includes, without limitation, electrical separation of one or more components from another component or from one another. As used herein, the term “circuit” is a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.

Various types of memory modules 10 are compatible with embodiments described herein. For example, memory modules 10 having memory capacities of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, as well as other capacities, are compatible with embodiments described herein. Certain embodiments described herein are applicable to various frequencies including, but not limited to 100 MHz, 200 MHz, 400 MHz, 800 MHz, and above. In addition, memory modules 10 having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the memory module 10 comprises a printed circuit board on which the memory devices 30 are mounted, a plurality of edge connectors configured to be electrically coupled to a corresponding plurality of contacts of a module slot of the computer system, and a plurality of electrical conduits which electrically couple the memory devices 30 to the circuit 40 and which electrically couple the circuit 40 to the edge connectors. Furthermore, memory modules 10 compatible with embodiments described herein include, but are not limited to, single in-line memory modules (SIMMs), dual in-line memory modules (DIMMs), small-outline DIMMs (SO-DIMMs), unbuffered DIMMs (UDIMMs), registered DIMMs (RDIMMs), fully-buffered DIMM (FBDIMM), rank-buffered DIMMs (RBDIMMs), mini-DIMMs, and micro-DIMMs.

Memory devices 30 compatible with embodiments described herein include, but are not limited to, random-access memory (RAM), dynamic random-access memory (DRAM), synchronous DRAM (SDRAM), and double-data-rate DRAM (e.g., SDR, DDR-1, DDR-2, DDR-3). In addition, memory devices 30 having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with embodiments described herein. Memory devices 30 compatible with embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBGA), micro-BGA (μ BGA), mini-BGA (mBGA), and chip-scale packaging (CSP). Memory devices 30 compatible with embodiments described herein are available from a number of sources, including but not limited to, Samsung Semiconductor, Inc. of San Jose, Calif., Infineon Technologies AG of San Jose, Calif., and Micron Technology, Inc. of Boise, Id. Persons skilled in the art can select appropriate memory devices 30 in accordance with certain embodiments described herein.

In certain embodiments, the plurality of memory devices 30 comprises a first number of memory devices 30. In certain such embodiments, the circuit 40 selectively isolates a second number of the memory devices 30 from the computer system, with the second number less than the first number.

In certain embodiments, the plurality of memory devices 30 are arranged in a first number of ranks. For example, in certain embodiments, the memory devices 30 are arranged in two ranks, as schematically illustrated by FIG. 1. In other

6

embodiments, the memory devices 30 are arranged in four ranks. Other numbers of ranks of the memory devices 30 are also compatible with embodiments described herein.

In certain embodiments, the circuit comprises a logic element selected from a group consisting of: a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a custom-designed semiconductor device, and a complex programmable-logic device (CPLD). In certain embodiments, the logic element of the circuit 40 is a custom device. Sources of logic elements compatible with embodiments described herein include, but are not limited to, Lattice Semiconductor Corporation of Hillsboro, Oreg., Altera Corporation of San Jose, Calif., and Xilinx Incorporated of San Jose, Calif. In certain embodiments, the logic element comprises various discrete electrical elements, while in certain other embodiments, the logic element comprises one or more integrated circuits.

In certain embodiments, the circuit 40 further comprises one or more switches which are operatively coupled to the logic element to receive control signals from the logic element. Examples of switches compatible with certain embodiments described herein include, but are not limited to, field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex.

FIG. 2 schematically illustrates a circuit diagram of two memory devices 30a, 30b of a conventional memory module showing the interconnections between the DQ data signal lines 102a, 102b of the memory devices 30a, 30b and the DQS data strobe signal lines 104a, 104b of the memory devices 30a, 30b. Each of the memory devices 30a, 30b has a plurality of DQ data signal lines and a plurality of DQS data strobe signal lines, however, for simplicity, FIG. 2 only illustrates a single DQ data signal line and a single DQS data strobe signal line for each memory device 30a, 30b. The DQ data signal lines 102a, 102b and the DQS data strobe signal lines 104a, 104b are typically conductive traces etched on the printed circuit board of the memory module. As shown in FIG. 2, each of the memory devices 30a, 30b has their DQ data signal lines 102a, 102b electrically coupled to a common DQ line 112 and their DQS data strobe signal lines 104a, 104b electrically coupled to a common DQS line 114. The common DQ line 112 and the common DQS line 114 are electrically coupled to the memory controller 20 of the computer system. Thus, the computer system is exposed to the loads of both memory devices 30a, 30b concurrently.

In certain embodiments, the circuit 40 selectively isolates the loads of at least some of the memory devices 30 from the computer system. The circuit 40 of certain embodiments is configured to present a significantly reduced load to the computer system. In certain embodiments in which the memory devices 30 are arranged in a plurality of ranks, the circuit 40 selectively isolates the loads of some (e.g., one or more) of the ranks of the memory module 10 from the computer system. In certain other embodiments, the circuit 40 selectively isolates the loads of all of the ranks of the memory module 10 from the computer system. For example, when a memory module 10 is not being accessed by the computer system, the capacitive load on the memory controller 20 of the computer system by the memory module 10 can be substantially reduced to the capacitive load of the circuit 40 of the memory module 10.

As schematically illustrated by FIGS. 3A and 3B, an example memory module 10 compatible with certain embodiments described herein comprises a circuit 40 which selectively isolates one or both of the DQ data signal lines 102a, 102b of the two memory devices 30a, 30b from the

US 7,881,150 B2

7

common DQ data signal line 112 coupled to the computer system. Thus, the circuit 40 selectively allows a DQ data signal to be transmitted from the memory controller 20 of the computer system to one or both of the DQ data signal lines 102a, 102b. In addition, the circuit 40 selectively allows one of a first DQ data signal from the DQ data signal line 102a of the first memory device 30a or a second DQ data signal from the DQ data signal line 102b of the second memory device 30b to be transmitted to the memory controller 20 via the common DQ data signal line 112 (see, e.g., triangles on the DQ and DQS lines of FIGS. 3A and 3B which point towards the memory controller). While various figures of the present application denote read operations by use of DQ and DQS lines which have triangles pointing towards the memory controller, certain embodiments described herein are also compatible with write operations (e.g., as would be denoted by triangles on the DQ or DQS lines pointing away from the memory controller).

For example, in certain embodiments, the circuit 40 comprises a pair of switches 120a, 120b on the DQ data signal lines 102a, 102b as schematically illustrated by FIG. 3A. Each switch 120a, 120b is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ signal line 112, the DQ data signal line 102b to the common DQ signal line 112, or both DQ data signal lines 102a, 102b to the common DQ signal line 112. In certain other embodiments, the circuit 40 comprises a switch 120 electrically coupled to both of the DQ data signal lines 102a, 102b, as schematically illustrated by FIG. 3B. The switch 120 is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ signal line 112, the DQ data signal line 102b to the common DQ signal line 112, or both DQ data signal lines 102a, 102b to the common DQ signal line 112. Circuits 40 having other configurations of switches are also compatible with embodiments described herein. While each of the memory devices 30a, 30b has a plurality of DQ data signal lines and a plurality of DQS data strobe signal lines, FIGS. 3A and 3B only illustrate a single DQ data signal line and a single DQS data strobe signal line for each memory device 30a, 30b for simplicity. The configurations schematically illustrated by FIGS. 3A and 3B can be applied to all of the DQ data signal lines and DQS data strobe signal lines of the memory module 10.

In certain embodiments, the circuit 40 selectively isolates the loads of ranks of memory devices 30 from the computer system. As schematically illustrated in FIGS. 4A and 4B, example memory modules 10 compatible with certain embodiments described herein comprise a first number of memory devices 30 arranged in a first number of ranks 32. The memory modules 10 of FIGS. 4A and 4B comprises two ranks 32a, 32b, with each rank 32a, 32b having a corresponding set of DQ data signal lines and a corresponding set of DQS data strobe lines. Other numbers of ranks (e.g., four ranks) of memory devices 30 of the memory module 10 are also compatible with certain embodiments described herein. For simplicity, FIGS. 4A and 4B illustrate only a single DQ data signal line and a single DQS data strobe signal line from each rank 32.

The circuit 40 of FIG. 4A selectively isolates one or more of the DQ data signal lines 102a, 102b of the two ranks 32a, 32b from the computer system. Thus, the circuit 40 selectively allows a DQ data signal to be transmitted from the memory controller 20 of the computer system to the memory devices 30 of one or both of the ranks 32a, 32b via the DQ data signal lines 102a, 102b. In addition, the circuit 40 selectively allows one of a first DQ data signal from the DQ data signal line 102a of the first rank 32a and a second DQ data signal

8

from the DQ data signal line 102b of the second rank 32b to be transmitted to the memory controller 20 via the common DQ data signal line 112. For example, in certain embodiments, the circuit 40 comprises a pair of switches 120a, 120b on the DQ data signal lines 102a, 102b as schematically illustrated by FIG. 4A. Each switch 120a, 120b is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ data signal line 112, the DQ data signal line 102b to the common DQ data signal line 112, or both DQ data signal lines 102a, 102b to the common DQ data signal line 112. In certain other embodiments, the circuit 40 comprises a switch 120 electrically coupled to both of the DQ data signal lines 102a, 102b, as schematically illustrated by FIG. 4B. The switch 120 is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ data signal line 112, the DQ data signal line 102b to the common DQ data signal line 112, or both DQ data signal lines 102a, 102b to the common DQ data signal line 112. Circuits 40 having other configurations of switches are also compatible with embodiments described herein.

In the example embodiments schematically illustrated by FIGS. 3A, 3B, 4A, and 4B, the circuit 40 comprises a logic element which is integral with and comprises the switches 120 which are coupled to the DQ data signal lines and the DQS data strobe signal lines. In certain such embodiments, each switch 120 comprises a data path multiplexer/demultiplexer. In certain other embodiments, the circuit 40 comprises a logic element 122 which is a separate component operatively coupled to the switches 120, as schematically illustrated by FIGS. 5A-5D. The one or more switches 120 are operatively coupled to the logic element 122 to receive control signals from the logic element 122 and to selectively electrically couple one or more data signal lines to a common data signal line. Example switches compatible with embodiments described herein include, but are not limited to, field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex. Example logic elements 122 compatible with certain embodiments described herein include, but are not limited to, programmable-logic devices (PLD), application-specific integrated circuits (ASIC), field-programmable gate arrays (FPGA), custom-designed semiconductor devices, and complex programmable-logic devices (CPLD). Example logic elements 122 are available from Lattice Semiconductor Corporation of Hillsboro, Oreg., Altera Corporation of San Jose, Calif., and Xilinx Incorporated of San Jose, Calif.

In certain embodiments, the load isolation provided by the circuit 40 advantageously allows the memory module 10 to present a reduced load (e.g., electrical load, such as capacitive load, inductive load, or impedance load) to the computer system by selectively switching between the two ranks of memory devices 30 to which it is coupled. This feature is used in certain embodiments in which the load of the memory module 10 may otherwise limit the number of ranks or the number of memory devices per memory module. In certain embodiments, the memory module 10 operates as having a data path rank buffer which advantageously isolates the ranks of memory devices 30 of the memory module 10 from one another, from the ranks on other memory modules, and from the computer system. This data path rank buffer of certain embodiments advantageously provides DQ-DQS paths for each rank or sets of ranks of memory devices which are separate from one another, or which are separate from the memory controller of the computer system. In certain embodiments, the load isolation advantageously diminishes the effects of capacitive loading, jitter and other sources of

US 7,881,150 B2

9

noise. In certain embodiments, the load isolation advantageously simplifies various other aspects of operation of the memory module 10, including but not limited to, setup-and-hold time, clock skew, package skew, and process, temperature, voltage, and transmission line variations.

For certain memory module applications that utilize multiple ranks of memory, increased load on the memory bus can degrade speed performance. In certain embodiments described herein, selectively isolating the loads of the ranks of memory devices 30 advantageously decreases the load on the computer system, thereby allowing the computer system (e.g., server) to run faster with improved signal integrity. In certain embodiments, load isolation advantageously provides system memory with reduced electrical loading, thereby improving the electrical topology to the memory controller 20. In certain such embodiments, the speed and the memory

10

density of the computer system are advantageously increased without sacrificing one for the other.

In certain embodiments, load isolation advantageously increases the size of the memory array supported by the 5 memory controller 20 of the computer system. The larger memory array has an increased number of memory devices 30 and ranks of memory devices 30 of the memory module 10, with a corresponding increased number of chip selects. Certain embodiments described herein advantageously provide 10 more system memory using fewer chip selects, thereby avoiding the chip select limitation of the memory controller.

An exemplary section of Verilog code corresponding to logic compatible with a circuit 40 which provides load isolation is listed below in Example 1. The exemplary code of 15 Example 1 corresponds to a circuit 40 comprising six FET switches for providing load isolation to DQ and DQS lines.

EXAMPLE 1

```

//===== declarations
reg rasN_R, casN_R, weN_R;
wire actv_cmd_R, pch_cmd_R, pch_all_cmd_R, ap_xfr_cmd_R_R;
wire xfr_cmd_R, mrs_cmd, rd_cmd_R;
/- ----- DDR 2 FET
reg brs0N_R;           // registered chip sel
reg brs1N_R;           // registered chip sel
reg brs2N_R;           // registered chip sel
reg brs3N_R;           // registered chip sel
wire sel;
wire sel_01;
wire sel_23;
wire rd_R1;
wire wr_cmd_R, wr_R1;
reg rd_R2, rd_R3, rd_R4, rd_R5;
reg wr_R2, wr_R3, wr_R4, wr_R5;
reg enfet1, enfet2, enfet3, enfet4, enfet5, enfet6;
wire wr_01_R1, wr_23_R1;
reg wr_01_R2, wr_01_R3, wr_01_R4;
reg wr_23_R2, wr_23_R3, wr_23_R4;
wire rodt0_a, rodt0_b;

//===== logic
always @ (posedge clk_in)
begin
    brs0N_R <= brs0_in_N;      // cs0
    brs1N_R <= brs1_in_N;      // cs1
    brs2N_R <= brs2_in_N;      // cs2
    brs3N_R <= brs3_in_N;      // cs3
    rasN_R <= bras_in_N;
    casN_R <= brcas_in_N;
    weN_R <= bwe_in_N;
end
assign sel = ~brs0N_R | ~brs1N_R | ~brs2N_R | ~brs3N_R ;
assign sel_01 = ~brs0N_R | ~brs1N_R ;
assign sel_23 = ~brs2N_R | ~brs3N_R ;
assign actv_cmd_R = !rasN_R & casN_R & weN_R;          // activate cmd
assign pch_cmd_R = !rasN_R & casN_R & !weN_R;          // pchg cmd
assign xfr_cmd_R = rasN_R & !casN_R;                   // xfr cmd
assign mrs_cmd = !rasN_R & !casN_R & !weN_R;          // md reg set cmd
assign rd_cmd_R = rasN_R & !casN_R & weN_R;          // read cmd
assign wr_cmd_R = rasN_R & !casN_R & !weN_R;          // write cmd
//-----
assign rd_R1 = sel & rd_cmd_R; // rd cmd cyc 1
assign wr_R1 = sel & wr_cmd_R; // wr cmd cyc 1
//-----
always @ (posedge clk_in)
begin
    rd_R2 <= rd_R1 ;
    rd_R3 <= rd_R2;
    rd_R4 <= rd_R3;
    rd_R5 <= rd_R4;
    rd0_o_R6 <= rd0_o_R5;
    wr_R2 <= wr_R1 ;
    wr_R3 <= wr_R2;
    wr_R4 <= wr_R3;
    wr_R5 <= wr_R4;

```

US 7,881,150 B2

11

12

-continued

```

        end
//-----
assign wr_01_R1 = sel_01 & wr_cmd_R; // wr cmd cyc 1 for cs 2 & cs3
assign wr_23_R1 = sel_23 & wr_cmd_R; // wr cmd cyc 1 for cs 2 & cs3
always @(posedge clk_in)
begin
    wr_01_R2 <= wr_01_R1 ;
    wr_01_R3 <= wr_01_R2;
    wr_01_R4 <= wr_01_R3;
    wr_23_R2 <= wr_23_R1 ;
    wr_23_R3 <= wr_23_R2;
    wr_23_R4 <= wr_23_R3;
end
assign rodt0_ab = (rodt0)           // odt cmd from sys
    | (wr_23_R1) // wr 1st cyc to other ranks (assume single dimm per channel)
    | (wr_23_R2) // wr 2nd cyc to other ranks (assume single dimm per channel)
    | (wr_23_R3) // wr 3rd cyc to other ranks (assume single dimm per channel)
    ;
assign rodt1_ab = (rodt1)           // odt cmd from sys
    | (wr_01_R1) // wr 1st cyc to other ranks (assume single dimm per channel)
    | (wr_01_R2) // wr 2nd cyc to other ranks (assume single dimm per channel)
    | (wr_01_R3) // wr 3rd cyc to other ranks (assume single dimm per channel)
    ;
//-----
always @(posedge clk_in)
begin
if(
    | (rd_R2)          // pre-am rd
    | (rd_R3)          // 1st cyc of rd brst (cl3)
    | (rd_R4)          // 2nd cyc of rd brst (cl3)
    | (wr_R1)          // pre-am wr
    | (wr_R2)          // wr brst 1st cyc
    | (wr_R3)          // wr brst 2nd cyc
) begin
    enfet1 <= 1'b1;      // enable fet
    enfet2 <= 1'b1;      // enable fet
    enfet3 <= 1'b1;      // enable fet
    enfet4 <= 1'b1;      // enable fet
    enfet5 <= 1'b1;      // enable fet
    enfet6 <= 1'b1;      // enable fet
end
else begin
    enfet1 <= 1'b0;      // disable fet
    enfet2 <= 1'b0;      // disable fet
    enfet3 <= 1'b0;      // disable fet
    enfet4 <= 1'b0;      // disable fet
    enfet5 <= 1'b0;      // disable fet
    enfet6 <= 1'b0;      // disable fet
end
end

```

Back-to-Back Adjacent Read Commands

Due to their source synchronous nature, DDR SDRAM (e.g., DDR1, DDR2, DDR3) memory devices operate with a data transfer protocol which surrounds each burst of data strobes with a pre-amble time interval and a post-amble time interval. The pre-amble time interval provides a timing window for the receiving memory device to enable its data capture circuitry when a known valid level is present on the strobe signal to avoid false triggers of the memory device's capture circuit. The post-amble time interval provides extra time after the last strobe for this data capture to facilitate good signal integrity. In certain embodiments, when the computer system accesses two consecutive bursts of data from the same memory device, termed herein as a "back-to-back adjacent read," the post-amble time interval of the first read command and the pre-amble time interval of the second read command are skipped by design protocol to increase read efficiency. FIG. 6A shows an exemplary timing diagram of this "gapless" read burst for a back-to-back adjacent read condition from one memory device.

In certain embodiments, when the second read command accesses data from a different memory device than does the first read command, there is at least one time interval (e.g., 50 clock cycle) inserted between the data strobes of the two memory devices. This inserted time interval allows both read data bursts to occur without the post-amble time interval of the first read data burst colliding or otherwise interfering with the pre-amble time interval of the second read data burst. In 55 certain embodiments, the memory controller of the computer system inserts an extra clock cycle between successive read commands issued to different memory devices, as shown in the exemplary timing diagram of FIG. 6B for successive read accesses from different memory devices.

60 In typical computer systems, the memory controller is informed of the memory boundaries between the ranks of memory of the memory module prior to issuing read commands to the memory module. Such memory controllers can insert wait time intervals or clock cycles to avoid collisions or 65 interference between back-to-back adjacent read commands which cross memory device boundaries, which are referred to herein as "BBARX."

US 7,881,150 B2

13

In certain embodiments described herein in which the number of ranks 32 of the memory module 10 is doubled or quadrupled, the circuit 40 generates a set of output address and command signals so that the selection decoding is transparent to the computer system. However, in certain such embodiments, there are memory device boundaries of which the computer system is unaware, so there are occasions in which BBARX occurs without the cognizance of the memory controller 20 of the computer system. As shown in FIG. 7, the last data strobe of memory device "a" collides with the preamble time interval of the data strobe of memory device "b," resulting in a "collision window."

FIGS. 8A-8D schematically illustrate circuit diagrams of example memory modules 10 comprising a circuit 40 which multiplexes the DQS data strobe signal lines 104a, 104b of two ranks 32a, 32b from one another in accordance with certain embodiments described herein. While the DQS data strobe signal lines 104a, 104b of FIGS. 8A-8D correspond to two ranks 32a, 32b of memory devices 30, in certain other embodiments, the circuit 40 multiplexes the DQS data strobe signal lines 104a, 104b corresponding to two individual memory devices 30a, 30b.

FIG. 8A schematically illustrates a circuit diagram of an exemplary memory module 10 comprising a circuit 40 in accordance with certain embodiments described herein. In certain embodiments, BBARX collisions are avoided by a mechanism which electrically isolates the DQS data strobe signal lines 104a, 104b from one another during the transition from the first read data burst of one rank 32a of memory devices 30 to the second read data burst of another rank 32b of memory devices 30.

In certain embodiments, as schematically illustrated by FIG. 8A, the circuit 40 comprises a first switch 130a electrically coupled to a first DQS data strobe signal line 104a of a first rank 32a of memory devices 30 and a second switch 130b electrically coupled to a second DQS data strobe signal line 104b of a second rank 32b of memory devices 30. In certain embodiments, the time for switching the first switch 130a and the second switch 130b is between the two read data bursts (e.g., after the last DQS data strobe of the read data burst of the first rank 32a and before the first DQS data strobe of the read data burst of the second rank 32b). During the read data burst for the first rank 32a, the first switch 130a is enabled. After the last DQS data strobe of the first rank 32a and before the first DQS data strobe of the second rank 32b, the first switch 130a is disabled and the second switch 130b is enabled.

As shown in FIG. 8A, each of the ranks 32a, 32b otherwise involved in a BBARX collision have their DQS data strobe signal lines 104a, 104b selectively electrically coupled to the common DQS line 114 through the circuit 40. The circuit 40 of certain embodiments multiplexes the DQS data strobe signal lines 104a, 104b of the two ranks 32a, 32b of memory devices 30 from one another to avoid a BBARX collision.

In certain embodiments, as schematically illustrated by FIG. 8B, the circuit 40 comprises a switch 130 which multiplexes the DQS data strobe signal lines 104a, 104b from one another. For example, the circuit 40 receives a DQS data strobe signal from the common DQS data strobe signal line 114 and selectively transmits the DQS data strobe signal to the first DQS data strobe signal line 104a, to the second DQS data strobe signal line 104b, or to both DQS data strobe signal lines 104a, 104b. As another example, the circuit 40 receives

14

a first DQS data strobe signal from the first rank 32a of memory devices 30 and a second DQS data strobe signal from a second rank 32b of memory devices 30 and selectively switches one of the first and second DQS data strobe signals to the common DQS data strobe signal line 114.

In certain embodiments, the circuit 40 also provides the load isolation described above in reference to FIGS. 1-5. For example, as schematically illustrated by FIG. 8C, the circuit 40 comprises both the switch 120 for the DQ data signal lines 102a, 102b and the switch 130 for the DQS data strobe signal lines 104a, 104b. While in certain embodiments, the switches 130 are integral with a logic element of the circuit 40, in certain other embodiments, the switches 130 are separate components which are operatively coupled to a logic element 122 of the circuit 40, as schematically illustrated by FIG. 8D. In certain such embodiments, the control and timing of the switch 130 is performed by the circuit 40 which is resident on the memory module 10. Example switches 130 compatible with embodiments described herein include, but are not limited to field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex., and multiplexers, such as the SN74AUC2G53 2:1 analog multiplexer/demultiplexer available from Texas Instruments, Inc. of Dallas, Tex.

The circuit 40 of certain embodiments controls the isolation of the DQS data strobe signal lines 104a, 104b by monitoring commands received by the memory module 10 from the computer system and producing "windows" of operation whereby the appropriate switches 130 are activated or deactivated to enable and disable the DQS data strobe signal lines 104a, 104b to mitigate BBARX collisions. In certain other embodiments, the circuit 40 monitors the commands received by the memory module 10 from the computer system and selectively activates or deactivates the switches 120 to enable and disable the DQ data signal lines 102a, 102b to reduce the load of the memory module 10 on the computer system. In still other embodiments, the circuit 40 performs both of these functions together.

Command Signal Translation

Most high-density memory modules are currently built with 512-Megabit ("512-Mb") memory devices wherein each memory device has a 64M×8-bit configuration. For example, a 1-Gigabyte ("1-GB") memory module with error checking capabilities can be fabricated using eighteen such 512-Mb memory devices. Alternatively, it can be economically advantageous to fabricate a 1-GB memory module using lower-density memory devices and doubling the number of memory devices used to produce the desired word width. For example, by fabricating a 1-GB memory module using thirty-six 256-Mb memory devices with 64M×4-bit configuration, the cost of the resulting 1-GB memory module can be reduced since the unit cost of each 256-Mb memory device is typically lower than one-half the unit cost of each 512-Mb memory device. The cost savings can be significant, even though twice as many 256-Mb memory devices are used in place of the 512-Mb memory devices. For example, by using pairs of 512-Mb memory devices rather than single 1-Gb memory devices, certain embodiments described herein reduce the cost of the memory module by a factor of up to approximately five.

US 7,881,150 B2

15

Market pricing factors for DRAM devices are such that higher-density DRAM devices (e.g., 1-Gb DRAM devices) are much more than twice the price of lower-density DRAM devices (e.g., 512-Mb DRAM devices). In other words, the price per bit ratio of the higher-density DRAM devices is greater than that of the lower-density DRAM devices. This pricing difference often lasts for months or even years after the introduction of the higher-density DRAM devices, until volume production factors reduce the costs of the newer higher-density DRAM devices. Thus, when the cost of a higher-density DRAM device is more than the cost of two lower-density DRAM devices, there is an economic incentive for utilizing pairs of the lower-density DRAM devices to replace individual higher-density DRAM devices.

FIG. 9A schematically illustrates an exemplary memory module 10 compatible with certain embodiments described herein. The memory module 10 is connectable to a memory controller 20 of a computer system (not shown). The memory module 10 comprises a printed circuit board 210 and a plurality of memory devices 30 coupled to the printed circuit board 210. The plurality of memory devices 30 has a first number of memory devices 30. The memory module 10 further comprises a circuit 40 coupled to the printed circuit board 210. The circuit 40 receives a set of input address and command signals from the computer system. The set of input address and command signals correspond to a second number of memory devices 30 smaller than the first number of memory devices 30. The circuit 40 generates a set of output address and command signals in response to the set of input address and command signals. The set of output address and command signals corresponds to the first number of memory devices 30.

In certain embodiments, as schematically illustrated in FIG. 9A, the memory module 10 further comprises a phase-lock loop device 220 coupled to the printed circuit board 210 and a register 230 coupled to the printed circuit board 210. In certain embodiments, the phase-lock loop device 220 and the register 230 are each mounted on the printed circuit board 210. In response to signals received from the computer system, the phase-lock loop device 220 transmits clock signals to the plurality of memory devices 30, the circuit 40, and the register 230. The register 230 receives and buffers a plurality of command signals and address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select signals), and transmits corresponding signals to the appropriate memory devices 30. In certain embodiments, the register 230 comprises a plurality of register devices. While the phase-lock loop device 220, the register 230, and the circuit 40 are described herein in certain embodiments as being separate components, in certain other embodiments, two or more of the phase-lock loop device 220, the register 230, and the circuit 40 are portions of a single component. Persons skilled in the art are able to select a phase-lock loop device 220 and a register 230 compatible with embodiments described herein.

In certain embodiments, the memory module 10 further comprises electrical components which are electrically coupled to one another and are surface-mounted or embedded on the printed circuit board 210. These electrical components can include, but are not limited to, electrical conduits, resistors, capacitors, inductors, and transistors. In certain embodiments,

16

ments, at least some of these electrical components are discrete, while in other certain embodiments, at least some of these electrical components are constituents of one or more integrated circuits.

5 In certain embodiments, the printed circuit board 210 is mountable in a module slot of the computer system. The printed circuit board 210 of certain such embodiments has a plurality of edge connections electrically coupled to corresponding contacts of the module slot and to the various components of the memory module 10, thereby providing electrical connections between the computer system and the components of the memory module 10.

In certain embodiments, the plurality of memory devices 15 30 are arranged in a first number of ranks 32. For example, in certain embodiments, the memory devices 30 are arranged in four ranks 32a, 32b, 32c, 32d, as schematically illustrated by FIG. 9A. In certain other embodiments, the memory devices 20 30 are arranged in two ranks 32a, 32b, as schematically illustrated by FIG. 9B. Other numbers of ranks 32 of the memory devices 30 are also compatible with embodiments described herein.

As schematically illustrated by FIGS. 9A and 9B, in certain 25 embodiments, the circuit 40 receives a set of input command signals (e.g., refresh, precharge) and address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select signals) from the memory controller 20 of the computer system. 30 In response to the set of input address and command signals, the circuit 40 generates a set of output address and command signals.

In certain embodiments, the set of output address and command signals corresponds to a first number of ranks in which the plurality of memory devices 30 of the memory module 10 are arranged, and the set of input address and command signals corresponds to a second number of ranks per memory module for which the computer system is configured. The 40 second number of ranks in certain embodiments is smaller than the first number of ranks. For example, in the exemplary embodiment as schematically illustrated by FIG. 9A, the first number of ranks is four while the second number of ranks is two. In the exemplary embodiment of FIG. 9B, the first number of ranks is two while the second number of ranks is one. Thus, in certain embodiments, even though the memory module 10 actually has the first number of ranks of memory devices 30, the memory module 10 simulates a virtual memory module by operating as having the second number of ranks of memory devices 30. In certain embodiments, the memory module 10 simulates a virtual memory module when the number of memory devices 30 of the memory module 10 is larger than the number of memory devices 30 per memory module for which the computer system is configured to utilize. In certain embodiments, the circuit 40 comprises logic (e.g., address decoding logic, command decoding logic) which translates between a system memory domain of the computer system and a physical memory domain of the memory module 10.

In certain embodiments, the computer system is configured for a number of ranks per memory module which is smaller than the number of ranks in which the memory devices 30 of the memory module 10 are arranged. In certain such embodiments, the computer system is configured for two ranks of

US 7,881,150 B2

17

memory per memory module (providing two chip-select signals CS₀, CS₁) and the plurality of memory modules 30 of the memory module 10 are arranged in four ranks, as schematically illustrated by FIG. 9A. In certain other such embodiments, the computer system is configured for one rank of memory per memory module (providing one chip-select signal CS₀) and the plurality of memory modules 30 of the memory module 10 are arranged in two ranks, as schematically illustrated by FIG. 9B.

In the exemplary embodiment schematically illustrated by FIG. 9A, the memory module 10 has four ranks of memory devices 30 and the computer system is configured for two ranks of memory devices per memory module. The memory module 10 receives row/column address signals or signal bits (A₀-A_{n+1}), bank address signals (BA₀-BA_m), chip-select signals (CS₀ and CS₁), and command signals (e.g., refresh, precharge, etc.) from the computer system. The A₀-A_n row/column address signals are received by the register 230, which buffers these address signals and sends these address signals to the appropriate ranks of memory devices 30. The circuit 40 receives the two chip-select signals (CS₀, CS₁) and one row/column address signal (A_{n+1}) from the computer system. Both the circuit 40 and the register 230 receive the bank address signals (BA₀-BA_m) and at least one command signal (e.g., refresh, precharge, etc.) from the computer system.

Logic Tables

Table 1 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices 30 using chip-select signals.

TABLE 1

State	CS ₀	CS ₁	A _{n+1}	Command	CS _{0A}	CS _{0B}	CS _{1A}	CS _{1B}
1	0	1	0	Active	0	1	1	1
2	0	1	1	Active	1	0	1	1
3	0	1	x	Active	0	0	1	1
4	1	0	0	Active	1	1	0	1
5	1	0	1	Active	1	1	1	0
6	1	0	x	Active	1	1	0	0
7	1	1	x	x	1	1	1	1

Note:

1. CS₀, CS₁, CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are active low signals.

2. A_{n+1} is an active high signal.

3. 'x' is a Don't Care condition.

4. Command involves a number of command signals that define operations such as refresh, precharge, and other operations.

30

25

35

40

45

50

55

60

65

70

75

80

85

90

95

100

18

In Logic State 1: CS₀ is active low, A_{n+1} is non-active, and Command is active. CS_{0A} is pulled low, thereby selecting Rank 0.

In Logic State 2: CS₀ is active low, A_{n+1} is active, and Command is active. CS_{0B} is pulled low, thereby selecting Rank 1.

In Logic State 3: CS₀ is active low, A_{n+1} is Don't Care, and Command is active high. CS_{0A} and CS_{0B} are pulled low, thereby selecting Ranks 0 and 1.

In Logic State 4: CS₁ is active low, A_{n+1} is non-active, and Command is active. CS_{1A} is pulled low, thereby selecting Rank 2.

In Logic State 5: CS₁ is active low, A_{n+1} is active, and Command is active. CS_{1B} is pulled low, thereby selecting Rank 3.

In Logic State 6: CS₁ is active low, A_{n+1} is Don't Care, and Command is active. CS_{1A} and CS_{1B} are pulled low, thereby selecting Ranks 2 and 3.

In Logic State 7: CS₀ and CS₁ are pulled non-active high, which deselects all ranks, i.e., CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are pulled high.

The "Command" column of Table 1 represents the various commands that a memory device (e.g., a DRAM device) can execute, examples of which include, but are not limited to, activation, read, write, precharge, and refresh. In certain embodiments, the command signal is passed through to the selected rank only (e.g., state 4 of Table 1). In such embodiments, the command signal (e.g., read) is sent to only one memory device or the other memory device so that data is supplied from one memory device at a time. In other embodiments, the command signal is passed through to both associated ranks (e.g., state 6 of Table 1). In such embodiments, the command signal (e.g., refresh) is sent to both memory devices to ensure that the memory content of the memory devices remains valid over time. Certain embodiments utilize a logic table such as that of Table 1 to simulate a single memory device from two memory devices by selecting two ranks concurrently.

Table 2 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices 30 using gated CAS signals.

TABLE 2

CS*	RAS*	CAS*	WE*	Density Bit	A ₁₀	Command	CAS0*	CAS1*
1	x	x	x	x	x	NOP	x	x
0	1	1	1	x	x	NOP	1	1
0	0	1	1	0	x	ACTIVATE	1	1
0	0	1	1	1	x	ACTIVATE	1	1
0	1	0	1	0	x	READ	0	1
0	1	0	1	1	x	READ	1	0
0	1	0	0	0	x	WRITE	0	1
0	1	0	0	1	x	WRITE	1	0
0	0	1	0	0	0	PRECHARGE	1	1
0	0	1	0	1	0	PRECHARGE	1	1
0	0	1	0	x	1	PRECHARGE	1	1
0	0	0	0	x	x	MODE REG SET	0	0
0	0	0	1	x	x	REFRESH	0	0

US 7,881,150 B2

19

In certain embodiments in which the density bit is a row address bit, for read/write commands, the density bit is the value latched during the activate command for the selected bank.

Serial-Presence-Detect Device

Memory modules typically include a serial-presence detect (SPD) device 240 (e.g., an electrically-erasable-programmable read-only memory or EEPROM device) comprising data which characterize various attributes of the memory module, including but not limited to, the number of row addresses the number of column addresses, the data width of the memory devices, the number of ranks, the memory density per rank, the number of memory devices, and the memory density per memory device. The SPD device 240 communicates this data to the basic input/output system (BIOS) of the computer system so that the computer system is informed of the memory capacity and the memory configuration available for use and can configure the memory controller properly for maximum reliability and performance.

For example, for a commercially-available 512-MB (64M×8-byte) memory module utilizing eight 512-Mb memory devices each with a 64M×8-bit configuration, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3: Defines the number of row address bits in the DRAM device in the memory module [13 for the 512-Mb memory device].

Byte 4: Defines the number of column address bits in the DRAM device in the memory module [11 for the 512-Mb memory device].

Byte 13: Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 512-Mb (64M×8-bit) memory device].

Byte 14: Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 512-Mb (64M×8-bit) memory device].

Byte 17: Defines the number of banks internal to the DRAM device used in the memory module [4 for the 512-Mb memory device].

In a further example, for a commercially-available 1-GB (128M×8-byte) memory module utilizing eight 1-Gb memory devices each with a 128M×8-bit configuration, as described above, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3: Defines the number of row address bits in the DRAM device in the memory module [14 for the 1-Gb memory device].

Byte 4: Defines the number of column address bits in the DRAM device in the memory module [11 for the 1-Gb memory device].

Byte 13: Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 1-Gb (128M×8-bit) memory device].

Byte 14: Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 1-Gb (128M×8-bit) memory device].

Byte 17: Defines the number of banks internal to the DRAM device used in the memory module [4 for the 1-Gb memory device].

In certain embodiments, the SPD device 240 comprises data which characterize the memory module 10 as having fewer ranks of memory devices than the memory module 10 actually has, with each of these ranks having more memory density. For example, for a memory module 10 compatible with certain embodiments described herein having two ranks of memory devices 30, the SPD device 240 comprises data

20

which characterizes the memory module 10 as having one rank of memory devices with twice the memory density per rank. Similarly, for a memory module 10 compatible with certain embodiments described herein having four ranks of

5 memory devices 30, the SPD device 240 comprises data which characterizes the memory module 10 as having two ranks of memory devices with twice the memory density per rank. In addition, in certain embodiments, the SPD device 240 comprises data which characterize the memory module 10 as having fewer memory devices than the memory module 10 actually has, with each of these memory devices having more memory density per memory device. For example, for a memory module 10 compatible with certain embodiments described herein, the SPD device 240 comprises data which 10 characterizes the memory module 10 as having one-half the number of memory devices that the memory module 10 actually has, with each of these memory devices having twice the memory density per memory device. Thus, in certain embodiments, the SPD device 240 informs the computer system of 15 the larger memory array by reporting a memory device density that is a multiple of the memory devices 30 resident on the memory module 10. Certain embodiments described herein advantageously do not require system level changes to hardware (e.g., the motherboard of the computer system) or to 20 software (e.g., the BIOS of the computer system).

FIG. 9C schematically illustrates an exemplary memory module 10 in accordance with certain embodiments described herein. The memory module 10 comprises a pair of substantially identical memory devices 31, 33. Each memory device 31, 33 has a first bit width, a first number of banks of memory locations, a first number of rows of memory locations, and a first number of columns of memory locations. The memory module 10 further comprises an SPD device 240 comprising data that characterizes the pair of memory devices 31, 33. The data characterize the pair of memory devices 31, 33 as a virtual memory device having a second bit width equal to twice the first bit width, a second number of banks of memory locations equal to the first number of banks, a second number of rows of memory locations equal to the first number 40 of rows, and a second number of columns of memory locations equal to the first number of columns.

In certain such embodiments, the SPD device 240 of the memory module 10 is programmed to describe the combined pair of lower-density memory devices 31, 33 as one virtual or 45 pseudo-higher-density memory device. In an exemplary embodiment, two 512-Mb memory devices, each with a 128M×4-bit configuration, are used to simulate one 1-Gb memory device having a 128M×8-bit configuration. The SPD device 240 of the memory module 10 is programmed to 50 describe the pair of 512-Mb memory devices as one virtual or pseudo-1-Gb memory device.

For example, to fabricate a 1-GB (128M×8-byte) memory module, sixteen 512-Mb (128M×4-bit) memory devices can be used. The sixteen 512-Mb (128M×4-bit) memory devices 55 are combined in eight pairs, with each pair serving as a virtual or pseudo-1-Gb (128M×8-bit) memory device. In certain such embodiments, the SPD device 240 contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3: 13 row address bits.

Byte 4: 12 column address bits.

Byte 13: 8 bits wide for the primary virtual 1-Gb (128M×8-bit) memory device.

Byte 14: 8 bits wide for the error checking virtual 1-Gb (128M×8-bit) memory device.

Byte 17: 4 banks.

In this exemplary embodiment, bytes 3, 4, and 17 are programmed to have the same values as they would have for

US 7,881,150 B2

21

a 512-MB (128M \times 4-byte) memory module utilizing 512-Mb (128M \times 4-bit) memory devices. However, bytes 13 and 14 of the SPD data are programmed to be equal to 8, corresponding to the bit width of the virtual or pseudo-higher-density 1-Gb (128M \times 8-bit) memory device, for a total capacity of 1-GB. Thus, the SPD data does not describe the actual-lower-density memory devices, but instead describes the virtual or pseudo-higher-density memory devices. The BIOS accesses the SPD data and recognizes the memory module as having 4 banks of memory locations arranged in 2^{13} rows and 2^{12} columns, with each memory location having a width of 8 bits rather than 4 bits.

In certain embodiments, when such a memory module 10 is inserted in a computer system, the computer system's memory controller then provides to the memory module 10 a set of input address and command signals which correspond to the number of ranks or the number of memory devices reported by the SPD device 240. For example, placing a two-rank memory module 10 compatible with certain embodiments described herein in a computer system compatible with one-rank memory modules, the SPD device 240 reports to the computer system that the memory module 10 only has one rank. The circuit 40 then receives a set of input address and command signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output address and command signals corresponding to two ranks to the appropriate memory devices 30 of the memory module 10.

Similarly, when a two-rank memory module 10 compatible with certain embodiments described herein is placed in a computer system compatible with either one- or two-rank memory modules, the SPD device 240 reports to the computer system that the memory module 10 only has one rank. The circuit 40 then receives a set of input address and command signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output address and command signals corresponding to two ranks to the appropriate memory devices 30 of the memory module 10.

Furthermore, a four-rank memory module 10 compatible with certain embodiments described herein simulates a two-rank memory module whether the memory module 10 is inserted in a computer system compatible with two-rank memory modules or with two- or four-rank memory modules. Thus, by placing a four-rank memory module 10 compatible with certain embodiments described herein in a module slot that is four-rank-ready, the computer system provides four chip-select signals, but the memory module 10 only uses two of the chip-select signals.

In certain embodiments, the circuit 40 comprises the SPD device 240 which reports the CAS latency (CL) to the memory controller of the computer system. The SPD device 240 of certain embodiments reports a CL which has one more cycle than does the actual operational CL of the memory array. In certain embodiments, data transfers between the memory controller and the memory module are registered for one additional clock cycle by the circuit 40. The additional clock cycle of certain embodiments is added to the transfer time budget with an incremental overall CAS latency. This extra cycle of time in certain embodiments advantageously provides sufficient time budget to add a buffer which electrically isolates the ranks of memory devices 30 from the memory controller 20. The buffer of certain embodiments comprises combinatorial logic, registers, and logic pipelines. In certain embodiments, the buffer adds a one-clock cycle time delay, which is equivalent to a registered DIMM, to accomplish the address decoding. The one-cycle time delay

22

of certain such embodiments provides sufficient time for read and write data transfers to provide the functions of the data path multiplexer/demultiplexer. Thus, for example, a DDR2 400-MHz memory system in accordance with embodiments described herein has an overall CAS latency of four, and uses memory devices with a CAS latency of three. In still other embodiments, the SPD device 240 does not utilize this extra cycle of time.

Memory Density Multiplication

In certain embodiments, two memory devices having a memory density are used to simulate a single memory device having twice the memory density, and an additional address signal bit is used to access the additional memory. Similarly, in certain embodiments, two ranks of memory devices having a memory density are used to simulate a single rank of memory devices having twice the memory density, and an additional address signal bit is used to access the additional memory. As used herein, such simulations of memory devices or ranks of memory devices are termed as "memory density multiplication," and the term "density transition bit" is used to refer to the additional address signal bit which is used to access the additional memory by selecting which rank of memory devices is enabled for a read or write transfer operation.

For example, for computer systems which are normally limited to using memory modules which have a single rank of 128M \times 4-bit memory devices, certain embodiments described herein enable the computer system to utilize memory modules which have double the memory (e.g., two ranks of 128M \times 4-bit memory devices). The circuit 40 of certain such embodiments provides the logic (e.g., command and address decoding logic) to double the number of chip selects, and the SPD device 240 reports a memory device density of 256M \times 4-bit to the computer system.

In certain embodiments utilizing memory density multiplication embodiments, the memory module 10 can have various types of memory devices 30 (e.g., DDR1, DDR2, DDR3, and beyond). The circuit 40 of certain such embodiments utilizes implied translation logic equations having variations depending on whether the density transition bit is a row, column, or internal bank address bit. In addition, the translation logic equations of certain embodiments vary depending on the type of memory module 10 (e.g., UDIMM, RDIMM, FBDIMM, etc.). Furthermore, in certain embodiments, the translation logic equations vary depending on whether the implementation multiplies memory devices per rank or multiplies the number of ranks per memory module.

Table 3A provides the numbers of rows and columns for DDR-1 memory devices, as specified by JEDEC standard JESD79D, "Double Data Rate (DDR) SDRAM Specification," published February 2004, and incorporated in its entirety by reference herein.

TABLE 3A

	128-Mb	256-Mb	512-Mb	1-Gb
Number of banks	4	4	4	4
Number of row address bits	12	13	13	14
Number of column address bits for "x 4" configuration	11	11	12	12
Number of column address bits for "x 8" configuration	10	10	11	11
Number of column address bits for "x 16" configuration	9	9	10	10

As described by Table 3A, 512-Mb (128M \times 4-bit) DRAM devices have 2^{13} rows and 2^{12} columns of memory locations,

US 7,881,150 B2

23

while 1-Gb (128M \times 8-bit) DRAM devices have 2^{14} rows and 2^{11} columns of memory locations. Because of the differences in the number of rows and the number of columns for the two types of memory devices, complex address translation procedures and structures would typically be needed to fabricate a 1-GB (128M \times 8-byte) memory module using sixteen 512-Mb (128M \times 4-bit) DRAM devices.

Table 3B shows the device configurations as a function of memory density for DDR2 memory devices.

TABLE 3B

	Number of Rows	Number of Columns	Number of Internal Banks	Page Size (x4s or x8s)
256 Mb	13	11	4	1 KB
512 Mb	14	11	4	1 KB
1 Gb	14	11	8	1 KB
2 Gb	15	11	8	1 KB
4 Gb	16	11	8	1 KB

Table 4 lists the corresponding density transition bit for the density transitions between the DDR2 memory densities of Table 3B.

TABLE 4

Density Transition	Density Transition Bit
256 Mb to 512 Mb	A ₁₃
512 Mb to 1 Gb	BA ₂
1 Gb to 2 Gb	A ₁₄
2 Gb to 4 Gb	A ₁₅

Other certain embodiments described herein utilize a transition bit to provide a transition from pairs of physical 4-Gb memory devices to simulated 8-Gb memory devices.

In an example embodiment, the memory module comprises one or more pairs of 256-Mb memory devices, with each pair simulating a single 512-Mb memory device. The simulated 512-Mb memory device has four internal banks while each of the two 256-Mb memory devices has four internal banks, for a total of eight internal banks for the pair of 256-Mb memory devices. In certain embodiments, the additional row address bit is translated by the circuit 40 to the rank selection between each of the two 256-Mb memory devices of the pair. Although there are eight total internal banks in the rank-converted memory array, the computer system is only aware of four internal banks. When the memory controller activates a row for a selected bank, the circuit 40 activates the same row for the same bank, but it does so for the selected rank according to the logic state of the additional row address bit A₁₃.

In another example embodiment, the memory module comprises one or more pairs of 512-Mb memory devices,

24

with each pair simulating a single 1-Gb memory device. The simulated 1-Gb memory device has eight internal banks while each of the two 512-Mb memory devices has four internal banks, for a total of eight internal banks for the pair of 512-Mb memory devices. In certain embodiments, the mapped BA₂ (bank 2) bit is used to select between the two ranks of 512-Mb memory devices to preserve the internal bank geometry expected by the memory controller of the computer system. The state of the BA₂ bit selects the upper or lower set of four banks, as well as the upper and lower 512-Mb rank.

In another example embodiment, the memory module comprises one or more pairs of 1-Gb memory devices, with each pair simulating a single 2-Gb memory device. Each of the two 1-Gb memory devices has eight internal banks for a total of sixteen internal banks, while the simulated 2-Gb memory device has eight internal banks. In certain embodiments, the additional row address bit translates to the rank selection between the two 1-Gb memory devices. Although there are sixteen total internal banks per pair of 1-Gb memory devices in the rank-converted memory array, the memory controller of the computer system is only aware of eight internal banks. When the memory controller activates a row of a selected bank, the circuit 40 activates the same row for the same bank, but it does so for the selected rank according to the logic state of the additional row address bit A₁₄.

The circuit 40 of certain embodiments provides substantially all of the translation logic used for the decoding (e.g., command and address decoding). In certain such embodiments, there is a fully transparent operational conversion from the “system memory” density domain of the computer system to the “physical memory” density domain of the memory module 10. In certain embodiments, the logic translation equations are programmed in the circuit 40 by hardware, while in certain other embodiments, the logic translation equations are programmed in the circuit 40 by software. Examples 1 and 2 provide exemplary sections of Verilog code compatible with certain embodiments described herein. As described more fully below, the code of Examples 1 and 2 includes logic to reduce potential problems due to “back-to-back adjacent read commands which cross memory device boundaries or “BBARX.” Persons skilled in the art are able to provide additional logic translation equations compatible with embodiments described herein.

An exemplary section of Verilog code compatible with memory density multiplication from 512 Mb to 1 Gb using DDR2 memory devices with the BA₂ density transition bit is listed below in Example 2. The exemplary code of Example 2 corresponds to a circuit 40 which receives one chip-select signal from the computer system and which generates two chip-select signals.

EXAMPLE 2

```

always @(posedge clk_in)
begin
    rs0N_R <= rs0_in_N; // cs0
    rasN_R <= ras_in_N;
    casN_R <= cas_in_N;
    weN_R <= we_in_N;
end
// Gated Chip Selects
assign   pcs0a_1 = (~rs0_in_N & ~ras_in_N & ~cas_in_N) // ref,md reg set
          | (~rs0_in_N & ras_in_N & cas_in_N) // ref exit, pwr dn
          | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in) // pchg all

```

US 7,881,150 B2

25

26

-continued

```

| (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & ~a10_in & ~ba2_in)// pchg single bnk
| (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N           & ~ba2_in)          // activate
| (~rs0_in_N & ras_in_N & ~cas_in_N           & ~ba2_in)          // xfr
;
assign   pcs0b_1 = (~rs0_in_N & ~ras_in_N & ~cas_in_N)           // ref,md reg set
| (~rs0_in_N & ras_in_N & cas_in_N)                      // ref exit, pwr dn
| (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in)      // pchg all
| (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N           & ~a10_in & ba2_in)// pchg single bnk
| (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N           & ba2_in)          // activate
| (~rs0_in_N & ras_in_N & ~cas_in_N           & ba2_in)          // xfr
;
//-----
always @ (posedge clk_in)
begin
    a4_r <= a4_in ;
    a5_r <= a5_in ;
    a6_r <= a6_in ;
    a10_r <= a10_in ;
    ba0_r <= ba0_in ;
    ba1_r <= ba1_in ;
    ba2_r <= ba2_in ;
    q_mrs_cmd_cyc1 <= q_mrs_cmd
end
///////////////////////////////
// determine the cas latency
/////////////////////////////
assign q_mrs_cmd_r = (!rasN_R & !casN_R & !weN_R
                      & !rs0N_R
                      & (!ba0_r & !ba1_r)
                      ;
                     // md reg set cmd
always @ (posedge clk_in)
if (~reset_N) // lmr
    cl3 <= 1'b1 ;
else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
    cl3 <= (~a6_r & a5_r & a4_r) ;
end
always @ (posedge clk_in)
if (~reset_N) // reset
    cl2 <= 1'b0 ;
else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
    cl2 <= (~a6_r & a5_r & ~a4_r) ;
end
always @ (posedge clk_in)
if (~reset_N) // reset
    cl4 <= 1'b0 ;
else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
    cl4 <= (a6_r & ~a5_r & ~a4_r) ;
end
always @ (posedge clk_in)
if (~reset_N) cl5 <= 1'b0 ;
else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
    cl5 <= (a6_r & ~a5_r & a4_r) ;
end
assign     pre_cyc2_enfet = (wr_cmd_cyc1 & acs_cyc1 & cl3)      // wr brst cl3 preamble
;
assign     pre_cyc3_enfet = (rd_cmd_cyc2 & cl3)                  // rd brst cl3 preamble
| (wr_cmd_cyc2 & cl3)                                         // wr brst cl3 1st pair
| (wr_cmd_cyc2 & cl4)                                         // wr brst cl4 preamble
;
assign     pre_cyc4_enfet = (wr_cmd_cyc3 & cl3)                  // wr brst cl3 2nd pair
| (wr_cmd_cyc3 & cl4)                                         // wr brst cl4 1st pair
| (rd_cmd_cyc3 & cl3)                                         // rd brst cl3 1st pair
| (rd_cmd_cyc3 & cl4)                                         // rd brst cl4 preamble
;
assign     pre_cyc5_enfet = (rd_cmd_cyc4 & cl3)                  // rd brst cl3 2nd pair
| (wr_cmd_cyc4 & cl4)                                         // wr brst cl4 2nd pair
| (rd_cmd_cyc4 & cl4)                                         // rd brst cl4 1st pair
;
// dq
assign     pre_dq_cyc = pre_cyc2_enfet
| pre_cyc3_enfet
| pre_cyc4_enfet
| pre_cyc5_enfet
;
assign     pre_dq_necy = enfet_cyc2

```

US 7,881,150 B2

27

28

-continued

```

| enfet_cyc3
| enfet_cyc4
| enfet_cyc5
;

// dqs
assign pre_dqsa_cyc = (pre_cyc2_enfet & ~ba2_r)
| (pre_cyc3_enfet & ~ba2_cyc2)
| (pre_cyc4_enfet & ~ba2_cyc3)
| (pre_cyc5_enfet & ~ba2_cyc4)
;
assign pre_dqsb_cyc = (pre_cyc2_enfet & ba2_r)
| (pre_cyc3_enfet & ba2_cyc2)
| (pre_cyc4_enfet & ba2_cyc3)
| (pre_cyc5_enfet & ba2_cyc4)
;
assign pre_dqsa_ncyc = (enfet_cyc2 & ~ba2_cyc2)
| (enfet_cyc3 & ~ba2_cyc3)
| (enfet_cyc4 & ~ba2_cyc4)
| (enfet_cyc5 & ~ba2_cyc5)
;
assign pre_dqsb_ncyc = (enfet_cyc2 & ba2_cyc2)
| (enfet_cyc3 & ba2_cyc3)
| (enfet_cyc4 & ba2_cyc4)
| (enfet_cyc5 & ba2_cyc5)
;
always @(posedge clk_in)
begin
    acs_cyc2 <= acs_cyc1 ;           // cs active
    ba2_cyc2 <= ba2_r ;
    ba2_cyc3 <= ba2_cyc2 ;
    ba2_cyc4 <= ba2_cyc3 ;
    ba2_cyc5 <= ba2_cyc4 ;
    rd_cmd_cyc2 <= rd_cmd_cyc1 & acs_cyc1;
    rd_cmd_cyc3 <= rd_cmd_cyc2 ;
    rd_cmd_cyc4 <= rd_cmd_cyc3 ;
    rd_cmd_cyc5 <= rd_cmd_cyc4 ;
    rd_cmd_cyc6 <= rd_cmd_cyc5 ;
    rd_cmd_cyc7 <= rd_cmd_cyc6 ;
    wr_cmd_cyc2 <= wr_cmd_cyc1 & acs_cyc1;
    wr_cmd_cyc3 <= wr_cmd_cyc2 ;
    wr_cmd_cyc4 <= wr_cmd_cyc3 ;
    wr_cmd_cyc5 <= wr_cmd_cyc4 ;
end
always @(negedge clk_in)
begin
    dq_ncyc <= dq_cyc;
    dqs_neyc_a <= dqs_cyc_a;
    dqs_neyc_b <= dqs_cyc_b;
end
// DQ FET enables
assign enq_fet1 = dq_cyc | dq_ncyc      ;
assign enq_fet2 = dq_cyc | dq_ncyc      ;
assign enq_fet3 = dq_cyc | dq_ncyc      ;
assign enq_fet4 = dq_cyc | dq_ncyc      ;
assign enq_fet5 = dq_cyc | dq_ncyc      ;
// DQS FET enables
assign ens_fet1a = dqs_cyc_a | dqs_ncyc_a ;
assign ens_fet2a = dqs_cyc_a | dqs_ncyc_a ;
assign ens_fet3a = dqs_cyc_a | dqs_ncyc_a ;
assign ens_fet1b = dqs_cyc_b | dqs_ncyc_b ;
assign ens_fet2b = dqs_cyc_b | dqs_ncyc_b ;
assign ens_fet3b = dqs_cyc_b | dqs_ncyc_b ;

```

Another exemplary section of Verilog code compatible with memory density multiplication from 256 Mb to 512 Mb using DDR2 memory devices and gated CAS signals with the row A₁₃ density transition bit is listed below in Example 3. The exemplary code of Example 3 corresponds to a circuit 40

⁵⁵ which receives one gated CAS signal from the computer system and which generates two gated CAS signals.

EXAMPLE 3

```

// latched a13 flags cs0, banks 0-3
always @(posedge clk_in)
    if (actv_cmd_R & ~rs0N_R & ~bnk1_R & ~bnk0_R) // activate

```

US 7,881,150 B2

29

30

-continued

```

begin
    1_a13_00 <= a13_r;
end
always @ (posedge clk_in)
if (actv_cmd_R & ~rs0N_R & ~bnk1_R & bnk0_R) // activate
begin
    1_a13_01 <= a13_r;
end
always @ (posedge clk_in)
if (actv_cmd_R & ~rs0N_R & bnk1_R & ~bnk0_R) // activate
begin
    1_a13_10 <= a13_r;
end
always @ (posedge clk_in)
if (actv_cmd_R & ~rs0N_R & bnk1_R & bnk0_R) // activate
begin
    1_a13_11 <= a13_r;
end
// gated cas
assign cas_i = ~(casN_R);
assign cas0_o = (~rasN_R & cas_i)
| (rasN_R & ~1_a13_00 & ~bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & ~1_a13_01 & ~bnk1_R & bnk0_R & cas_i)
| (rasN_R & ~1_a13_10 & bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & ~1_a13_11 & bnk1_R & bnk0_R & cas_i)
;
assign cas1_o = (~rasN_R & cas_i)
| (rasN_R & 1_a13_00 & ~bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & 1_a13_01 & ~bnk1_R & bnk0_R & cas_i)
| (rasN_R & 1_a13_10 & bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & 1_a13_11 & bnk1_R & bnk0_R & cas_i)
;
assign pcas_0_N = ~cas0_o;
assign pcas_1_N = ~cas1_o;
assign rd0_o_R1 = rasN_R & cas0_o & weN_R & ~rs0N_R; // rnk0 rd cmd cyc
assign rd1_o_R1 = rasN_R & cas1_o & weN_R & ~rs0N_R; // rnk1 rd cmd cyc
assign wr0_o_R1 = rasN_R & cas0_o & ~weN_R & ~rs0N_R; // rnk0 wr cmd cyc
assign wr1_o_R1 = rasN_R & cas1_o & ~weN_R & ~rs0N_R; // rnk1 wr cmd cyc
always @ (posedge clk_in)
begin
    rd0_o_R2 <= rd0_o_R1;
    rd0_o_R3 <= rd0_o_R2;
    rd0_o_R4 <= rd0_o_R3;
    rd0_o_R5 <= rd0_o_R4;
    rd1_o_R2 <= rd1_o_R1;
    rd1_o_R3 <= rd1_o_R2;
    rd1_o_R4 <= rd1_o_R3;
    rd1_o_R5 <= rd1_o_R4;
    wr0_o_R2 <= wr0_o_R1;
    wr0_o_R3 <= wr0_o_R2;
    wr0_o_R4 <= wr0_o_R3;
    wr1_o_R2 <= wr1_o_R1;
    wr1_o_R3 <= wr1_o_R2;
    wr1_o_R4 <= wr1_o_R3;
end
always @ (posedge clk_in)
begin
if (
    (rd0_o_R2 & ~rd1_o_R4) // pre-am rd if no ped on rnk 1
    | rd0_o_R3 // 1st cyc of rd brst
    | rd0_o_R4 // 2nd cyc of rd brst
    | (rd0_o_R5 & ~rd1_o_R2 & ~rd1_o_R3) // post-rd cyc if no ped on rnk 1
    | (wr0_o_R1) // pre-am wr
    | wr0_o_R2 | wr0_o_R3 // wr brst 1st & 2nd cyc
    | (wr0_o_R4) // post-wr cyc (chgef9)
    | wr1_o_R1 | wr1_o_R2 | wr1_o_R3 | wr1_o_R4 // rank 1 (chgef9)
)
    en_fet_a <= 1'b1; // enable fet
else
    en_fet_a <= 1'b0; // disable fet
end
always @ (posedge clk_in)
begin
if (
    (rd1_o_R2 & ~rd0_o_R4)
    | rd1_o_R3
    | rd1_o_R4
    | (rd1_o_R5 & ~rd0_o_R2 & ~rd0_o_R3) // (chgef8)
    | (wr1_o_R1)
)

```

US 7,881,150 B2

31

32

-continued

```

| wr1_o_R2 | wr1_o_R3
| (wr1_o_R4)
| wr0_o_R1 | wr0_o_R2 | wr0_o_R3 | wr0_o_R4      // post-wr cyc      (chgef9)
      )
      en_fet_b <= 1'b1;                         //
else          en_fet_b <= 1'b0;
end

```

In certain embodiments, the chipset memory controller of the computer system uses the inherent behavioral characteristics of the memory devices (e.g., DDR2 memory devices) to optimize throughput of the memory system. For example, for each internal bank in the memory array, a row (e.g., 1 KB page) is advantageously held activated for an extended period of time. The memory controller, by anticipating a high number of memory accesses or hits to a particular region of memory, can exercise this feature to advantageously eliminate time-consuming pre-charge cycles. In certain such embodiments in which two half-density memory devices are transparently substituted for a single full-density memory device (as reported by the SPD device 240 to the memory controller), the memory devices advantageously support the “open row” feature.

FIG. 10A schematically illustrates an exemplary memory module 10 which doubles the rank density in accordance with certain embodiments described herein. The memory module 10 has a first memory capacity. The memory module 10 comprises a plurality of substantially identical memory devices 30 configured as a first rank 32a and a second rank 32b. In certain embodiments, the memory devices 30 of the first rank 32a are configured in pairs, and the memory devices 30 of the second rank 32b are also configured in pairs. In certain embodiments, the memory devices 30 of the first rank 32a are configured with their respective DQS pins tied together and the memory devices 30 of the second rank 32b are configured with their respective DQS pins tied together, as described more fully below. The memory module 10 further comprises a circuit 40 which receives a first set of address and command signals from a memory controller (not shown) of the computer system. The first set of address and command signals is compatible with a second memory capacity substantially equal to one-half of the first memory capacity. The circuit 40 translates the first set of address and command signals into a second set of address and command signals which is compatible with the first memory capacity of the memory module 10 and which is transmitted to the first rank 32a and the second rank 32b.

The first rank 32a of FIG. 10A has 18 memory devices 30 and the second rank 32b of FIG. 10A has 18 memory devices 30. Other numbers of memory devices 30 in each of the ranks 32a, 32b are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. 10A, the memory module 10 has a width of 8 bytes (or 64 bits) and each of the memory devices 30 of FIG. 10A has a bit width of 4 bits. The 4-bit-wide (“×4”) memory devices 30 of FIG. 10A have one-half the width, but twice the depth of 8-bit-wide (“×8”) memory devices. Thus, each pair of “×4” memory devices 30 has the same density as a single “×8” memory device, and pairs of “×4” memory devices 30 can be used instead of individual “×8” memory devices to provide the memory density of the memory module 10. For example, a

pair of 512-Mb 128M×4-bit memory devices has the same memory density as a 1-Gb 128M×8-bit memory device.

For two “×4” memory devices 30 to work in tandem to mimic a “×8” memory device, the relative DQS pins of the two memory devices 30 in certain embodiments are advantageously tied together, as described more fully below. In addition, to access the memory density of a high-density memory module 10 comprising pairs of “×4” memory devices 30, an additional address line is used. While a high-density memory module comprising individual “×8” memory devices with the next-higher density would also utilize an additional address line, the additional address lines are different in the two memory module configurations.

For example, a 1-Gb 128M×8-bit DRAM memory device uses row addresses A₁₃-A₀ and column addresses A₁₁ and A₉-A₀. A pair of 512-Mb 128M×4-bit DDR-1 DRAM memory devices uses row addresses A₁₂-A₀ and column addresses A₁₂, A₁₁, and A₉-A₀. In certain embodiments, a memory controller of a computer system utilizing a 1-GB 128M×8 memory module 10 comprising pairs of the 512-Mb 128M×4 memory devices 30 supplies the address and command signals including the extra row address (A₁₃) to the memory module 10. The circuit 40 receives the address and command signals from the memory controller and converts the extra row address (A₁₃) into an extra column address (A₁₂).

FIG. 10B schematically illustrates an exemplary circuit 40 compatible with embodiments described herein. The circuit 40 is used for a memory module 10 comprising pairs of “×4” memory devices 30 which mimic individual “×8” memory devices. In certain embodiments, each pair has the respective DQS pins of the memory devices 30 tied together. In certain embodiments, as schematically illustrated by FIG. 10B, the circuit 40 comprises a programmable-logic device (PLD) 42, a first multiplexer 44 electrically coupled to the first rank 32a of memory devices 30, and a second multiplexer 46 electrically coupled to the second rank 32b of memory devices 30. In certain embodiments, the PLD 42 and the first and second multiplexers 44, 46 are discrete elements, while in other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can select an appropriate PLD 42, first multiplexer 44, and second multiplexer 46 in accordance with embodiments described herein.

In the exemplary circuit 40 of FIG. 10B, during a row access procedure (CAS is high), the first multiplexer 44 passes the A₁₂ address through to the first rank 32a, the second multiplexer 46 passes the A₁₂ address through to the second rank 32b, and the PLD 42 saves or latches the A₁₃ address from the memory controller. In certain embodiments, a copy of the A₁₃ address is saved by the PLD 42 for each of the internal banks (e.g., 4 internal banks) per memory device 30. During a subsequent column access procedure (CAS is low), the first multiplexer 44 passes the previously-saved A₁₃ address through to the first rank 32a as the A₁₂ address and the second multiplexer 46 passes the previously-saved A₁₃ address

US 7,881,150 B2

33

through to the second rank 32b as the A₁₂ address. The first rank 32a and the second rank 32b thus interpret the previously-saved A₁₃ row address as the current A₁₂ column address. In this way, in certain embodiments, the circuit 40 translates the extra row address into an extra column address in accordance with certain embodiments described herein.

Thus, by allowing two lower-density memory devices to be used rather than one higher-density memory device, certain embodiments described herein provide the advantage of using lower-cost, lower-density memory devices to build “next-generation” higher-density memory modules. Certain embodiments advantageously allow the use of lower-cost readily-available 512-Mb DDR-2 SDRAM devices to replace more expensive 1-Gb DDR-2 SDRAM devices. Certain embodiments advantageously reduce the total cost of the resultant memory module.

FIG. 11A schematically illustrates an exemplary memory module 10 which doubles number of ranks in accordance with certain embodiments described herein. The memory module 10 has a first plurality of memory locations with a first memory density. The memory module 10 comprises a plurality of substantially identical memory devices 30 configured as a first rank 32a, a second rank 32b, a third rank 32c, and a fourth rank 32d. The memory module 10 further comprises a circuit 40 which receives a first set of address and command signals from a memory controller (not shown). The first set of address and command signals is compatible with a second plurality of memory locations having a second memory density. The second memory density is substantially equal to one-half of the first memory density. The circuit 40 translates the first set of address and command signals into a second set of address and command signals which is compatible with the first plurality of memory locations of the memory module 10 and which is transmitted to the first rank 32a, the second rank 32b, the third rank 32c, and the fourth rank 32d.

Each rank 32a, 32b, 32c, 32d of FIG. 11A has 9 memory devices 30. Other numbers of memory devices 30 in each of the ranks 32a, 32b, 32c, 32d are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. 11A, the memory module 10 has a width of 8 bytes (or 64 bits) and each of the memory devices 30 of FIG. 11A has a bit width of 8 bits. Because the memory module 10 has twice the number of 8-bit-wide (“x8”) memory devices 30 as does a standard 8-byte-wide memory module, the memory module 10 has twice the density as does a standard 8-byte-wide memory module. For example, a 1-GB 128Mx8-bit memory module with 36 512-Mb 128Mx8-bit memory devices (arranged in four ranks) has twice the memory density as a 512-Mb 128Mx8-bit memory module with 18 512-Mb 128Mx8-bit memory devices (arranged in two ranks).

To access the additional memory density of the high-density memory module 10, the two chip-select signals (CS₀, CS₁) are used with other address and command signals to gate a set of four gated CAS signals. For example, to access the additional ranks of four-rank 1-GB 128Mx8-byte DDR-1 DRAM memory module, the CS₀ and CS₁ signals along with the other address and command signals are used to gate the CAS signal appropriately, as schematically illustrated by FIG. 11A. FIG. 11B schematically illustrates an exemplary circuit 40 compatible with embodiments described herein. In certain embodiments, the circuit 40 comprises a programmable-logic device (PLD) 42 and four “OR” logic elements 52, 54, 56, 58 electrically coupled to corresponding ranks 32a, 32b, 32c, 32d of memory devices 30.

In certain embodiments, the PLD 42 comprises an ASIC, an FPGA, a custom-designed semiconductor device, or a

34

CPLD. In certain embodiments, the PLD 42 and the four “OR” logic elements 52, 54, 56, 58 are discrete elements, while in other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can select an appropriate PLD 42 and appropriate “OR” logic elements 52, 54, 56, 58 in accordance with embodiments described herein.

In the embodiment schematically illustrated by FIG. 11B, the PLD 42 transmits each of the four “enabled CAS” (ENCAS_{0a}, ENCAS_{0b}, ENCAS_{1a}, ENCAS_{1b}) signals to a corresponding one of the “OR” logic elements 52, 54, 56, 58. The CAS signal is also transmitted to each of the four “OR” logic elements 52, 54, 56, 58. The CAS signal and the “enabled CAS” signals are “low” true signals. By selectively activating each of the four “enabled CAS” signals which are inputted into the four “OR” logic elements 52, 54, 56, 58, the PLD 42 is able to select which of the four ranks 32a, 32b, 32c, 32d is active.

In certain embodiments, the PLD 42 uses sequential and combinatorial logic procedures to produce the gated CAS signals which are each transmitted to a corresponding one of the four ranks 32a, 32b, 32c, 32d. In certain other embodiments, the PLD 42 instead uses sequential and combinatorial logic procedures to produce four gated chip-select signals (e.g., CS_{0a}, CS_{0b}, CS_{1a}, and CS_{1b}) which are each transmitted to a corresponding one of the four ranks 32a, 32b, 32c, 32d.

Tied Data Strobe Signal Pins

For proper operation, the computer system advantageously recognizes a 1-GB memory module comprising 256-Mb memory devices with 64Mx4-bit configuration as a 1-GB memory module having 512-Mb memory devices with 64Mx8-bit configuration (e.g., as a 1-GB memory module with 128Mx8-byte configuration). This advantageous result is desirably achieved in certain embodiments by electrically connecting together two output signal pins (e.g., DQS or data strobe pins) of the two 256-Mb memory devices such that both output signal pins are concurrently active when the two memory devices are concurrently enabled. The DQS or data strobe is a bi-directional signal that is used during both read cycles and write cycles to validate or latch data. As used herein, the terms “tying together” or “tied together” refer to a configuration in which corresponding pins (e.g., DQS pins) of two memory devices are electrically connected together and are concurrently active when the two memory devices are concurrently enabled (e.g., by a common chip-select or CS signal). Such a configuration is different from standard memory module configurations in which the output signal pins (e.g., DQS pins) of two memory devices are electrically coupled to the same source, but these pins are not concurrently active since the memory devices are not concurrently enabled. However, a general guideline of memory module design warns against tying together two output signal pins in this way.

FIGS. 12 and 13 schematically illustrate a problem which may arise from tying together two output signal pins. FIG. 12 schematically illustrates an exemplary memory module 305 in which a first DQS pin 312 of a first memory device 310 is electrically connected to a second DQS pin 322 of a second memory device 320. The two DQS pins 312, 322 are both electrically connected to a memory controller 330.

FIG. 13 is an exemplary timing diagram of the voltages applied to the two DQS pins 312, 322 due to non-simultaneous switching. As illustrated by FIG. 13, at time t₁, both the first DQS pin 312 and the second DQS pin 322 are high, so no current flows between them. Similarly, at time t₄, both the first

US 7,881,150 B2

35

DQS pin **312** and the second DQS pin **322** are low, so no current flows between them. However, for times between approximately t_2 and approximately t_3 , the first DQS pin **312** is low while the second DQS pin **322** is high. Under such conditions, a current will flow between the two DQS pins **312**, **322**. This condition in which one DQS pin is low while the other DQS pin is high can occur for fractions of a second (e.g., 0.8 nanoseconds) during the dynamic random-access memory (DRAM) read cycle. During such conditions, the current flowing between the two DQS pins **312**, **322** can be substantial, resulting in heating of the memory devices **310**, **320**, and contributing to the degradation of reliability and eventual failure of these memory devices.

A second problem may also arise from tying together two output signal pins. FIG. 14 schematically illustrates another exemplary memory module **305** in which a first DQS pin **312** of a first memory device **310** is electrically connected to a second DQS pin **322** of a second memory device **320**. The two DQS pins **312**, **322** of FIG. 14 are both electrically connected to a memory controller (not shown). The DQ (data input/output) pin **314** of the first memory device **310** and the corresponding DQ pin **324** of the second memory device **320** are each electrically connected to the memory controller by the DQ bus (not shown). Typically, each memory device **310**, **320** will have a plurality of DQ pins (e.g., eight DQ pins per memory device), but for simplicity, FIG. 14 only shows one DQ pin for each memory device **310**, **320**.

Each of the memory devices **310**, **320** of FIG. 14 utilizes a respective on-die termination or “ODT” circuit **332**, **334** which has termination resistors (e.g., 75 ohms) internal to the memory devices **310**, **320** to provide signal termination. Each memory device **310**, **320** has a corresponding ODT signal pin **362**, **364** which is electrically connected to the memory controller via an ODT bus **340**. The ODT signal pin **362** of the first memory device **310** receives a signal from the ODT bus **340** and provides the signal to the ODT circuit **332** of the first memory device **310**. The ODT circuit **332** responds to the signal by selectively enabling or disabling the internal termination resistors **352**, **356** of the first memory device **310**. This behavior is shown schematically in FIG. 14 by the switches **342**, **344** which are either closed (dash-dot line) or opened (solid line). The ODT signal pin **364** of the second memory device **320** receives a signal from the ODT bus **340** and provides the signal to the ODT circuit **334** of the second memory device **320**. The ODT circuit **334** responds to the signal by selectively enabling or disabling the internal termination resistors **354**, **358** of the second memory device **320**. This behavior is shown schematically in FIG. 14 by the switches **346**, **348** which are either closed (dash-dot line) or opened (solid line). The switches **342**, **344**, **346**, **348** of FIG. 14 are schematic representations of the operation of the ODT circuits **332**, **334**, and do not signify that the ODT circuits **332**, **334** necessarily include mechanical switches.

Examples of memory devices **310**, **320** which include such ODT circuits **332**, **334** include, but are not limited to, DDR2 memory devices. Such memory devices are configured to selectively enable or disable the termination of the memory device in this way in response to signals applied to the ODT signal pin of the memory device. For example, when the ODT signal pin **362** of the first memory device **310** is pulled high, the termination resistors **352**, **356** of the first memory device **310** are enabled. When the ODT signal pin **362** of the first memory device **310** is pulled low (e.g., grounded), the termination resistors **352**, **356** of the first memory device **310** are disabled. By selectively disabling the termination resistors of an active memory device, while leaving the termination resistors of inactive memory devices enabled, such configurations

36

advantageously preserve signal strength on the active memory device while continuing to eliminate signal reflections at the bus-die interface of the inactive memory devices.

In certain configurations, as schematically illustrated by FIG. 14, the DQS pins **312**, **322** of each memory device **310**, **320** are selectively connected to a voltage VTT through a corresponding termination resistor **352**, **354** internal to the corresponding memory device **310**, **320**. Similarly, in certain configurations, as schematically illustrated by FIG. 14, the DQ pins **314**, **324** are selectively connected to a voltage VTT through a corresponding termination resistor **356**, **358** internal to the corresponding memory device **310**, **320**. In certain configurations, rather than being connected to a voltage VTT, the DQ pins **314**, **324** and/or the DQS pins **312**, **322** are selectively connected to ground through the corresponding termination resistors **352**, **354**, **356**, **358**. The resistances of the internal termination resistors **352**, **354**, **356**, **358** are selected to clamp the voltages so as to reduce the signal reflections from the corresponding pins. In the configuration schematically illustrated by FIG. 14, each internal termination resistor **352**, **354**, **356**, **358** has a resistance of approximately 75 ohms.

When connecting the first memory device **310** and the second memory device **320** together to form a double word width, both the first memory device **310** and the second memory device **320** are enabled at the same time (e.g., by a common CS signal). Connecting the first memory device **310** and the second memory device **320** by tying the DQS pins **312**, **322** together, as shown in FIG. 14, results in a reduced effective termination resistance for the DQS pins **312**, **322**. For example, for the exemplary configuration of FIG. 14, the effective termination resistance for the DQS pins **312**, **322** is approximately 37.5 ohms, which is one-half the desired ODT resistance (for 75-ohm internal termination resistors) to reduce signal reflections since the internal termination resistors **352**, **354** of the two memory devices **310**, **320** are connected in parallel. This reduction in the termination resistance can result in signal reflections causing the memory device to malfunction.

FIG. 15 schematically illustrates an exemplary memory module **400** in accordance with certain embodiments described herein. The memory module **400** comprises a first memory device **410** having a first data strobe (DQS) pin **412** and a second memory device **420** having a second data strobe (DQS) pin **422**. The memory module **400** further comprises a first resistor **430** electrically coupled to the first DQS pin **412**. The memory module **400** further comprises a second resistor **440** electrically coupled to the second DQS pin **422** and to the first resistor **430**. The first DQS pin **412** is electrically coupled to the second DQS pin **422** through the first resistor **430** and through the second resistor **440**.

In certain embodiments, the memory module **400** is a 1-GB unbuffered Double Data Rate (DDR) Synchronous Dynamic RAM (SDRAM) high-density dual in-line memory module (DIMM). FIGS. 16A and 16B schematically illustrate a first side **462** and a second side **464**, respectively, of such a memory module **400** with eighteen 64Mx4-bit, DDR-1 SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board (PCB) **460**. In certain embodiments, the memory module **400** further comprises a phase-lock-loop (PLL) clock driver **470**, an EEPROM for serial-presence detect (SPD) data **480**, and decoupling capacitors (not shown) mounted on the PCB in parallel to suppress switching noise on VDD and VDDQ power supply for DDR-1 SDRAM. By using synchronous design, such memory modules **400** allow precise control of data transfer between the memory module **400** and the system controller.

US 7,881,150 B2

37

Data transfer can take place on both edges of the DQS signal at various operating frequencies and programming latencies. Therefore, certain such memory modules 400 are suitable for a variety of high-performance system applications.

In certain embodiments, the memory module 400 comprises a plurality of memory devices configured in pairs, each pair having a first memory device 410 and a second memory device 420. For example, in certain embodiments, a 128M \times 72-bit DDR SDRAM high-density memory module 400 comprises thirty-six 64M \times 4-bit DDR-1 SDRAM integrated circuits in FBGA packages configured in eighteen pairs. The first memory device 410 of each pair has the first DQS pin 412 electrically coupled to the second DQS pin 422 of the second memory device 420 of the pair. In addition, the first DQS pin 412 and the second DQS pin 422 are concurrently active when the first memory device 410 and the second memory device 420 are concurrently enabled.

In certain embodiments, the first resistor 430 and the second resistor 440 each has a resistance advantageously selected to reduce the current flow between the first DQS pin 412 and the second DQS pin 422 while allowing signals to propagate between the memory controller and the DQS pins 412, 422. In certain embodiments, each of the first resistor 430 and the second resistor 440 has a resistance in a range between approximately 5 ohms and approximately 50 ohms. For example, in certain embodiments, each of the first resistor 430 and the second resistor 440 has a resistance of approximately 22 ohms. Other resistance values for the first resistor 430 and the second resistor 440 are also compatible with embodiments described herein. In certain embodiments, the first resistor 430 comprises a single resistor, while in other embodiments, the first resistor 430 comprises a plurality of resistors electrically coupled together in series and/or in parallel. Similarly, in certain embodiments, the second resistor 440 comprises a single resistor, while in other embodiments, the second resistor 440 comprises a plurality of resistors electrically coupled together in series and/or in parallel.

FIGS. 17A and 17B schematically illustrate an exemplary embodiment of a memory module 400 in which the first resistor 430 and the second resistor 440 are used to reduce the current flow between the first DQS pin 412 and the second DQS pin 422. As schematically illustrated by FIG. 17A, the memory module 400 is part of a computer system 500 having a memory controller 510. The first resistor 430 has a resistance of approximately 22 ohms and the second resistor 440 has a resistance of approximately 22 ohms. The first resistor 430 and the second resistor 440 are electrically coupled in parallel to the memory controller 510 through a signal line 520 having a resistance of approximately 25 ohms. The first resistor 430 and the second resistor 440 are also electrically coupled in parallel to a source of a fixed termination voltage (identified by VTT in FIGS. 17A and 17B) by a signal line 540 having a resistance of approximately 47 ohms. Such an embodiment can advantageously be used to allow two memory devices having lower bit widths (e.g., 4-bit) to behave as a single virtual memory device having a higher bit width (e.g., 8-bit).

FIG. 17B schematically illustrates exemplary current-limiting resistors 430, 440 in conjunction with the impedances of the memory devices 410, 420. During an exemplary portion of a data read operation, the memory controller 510 is in a high-impedance condition, the first memory device 410 drives the first DQS pin 412 high (e.g., 2.7 volts), and the second memory device 420 drives the second DQS pin 422 low (e.g., 0 volts). The amount of time for which this condition occurs is approximated by the time between t₂ and t₃ of FIG. 13, which in certain embodiments is approximately

38

twice the tDQSQ (data strobe edge to output data edge skew time, e.g., approximately 0.8 nanoseconds). At least a portion of this time in certain embodiments is caused by simultaneous switching output (SSO) effects.

5 In certain embodiments, as schematically illustrated by FIG. 17B, the DQS driver of the first memory device 410 has a driver impedance R₁ of approximately 17 ohms, and the DQS driver of the second memory device 420 has a driver impedance R₄ of approximately 17 ohms. Because the upper network of the first memory device 410 and the first resistor 430 (with a resistance R₂ of approximately 22 ohms) is approximately equal to the lower network of the second memory device 420 and the second resistor 440 (with a resistance R₃ of approximately 22 ohms), the voltage at the midpoint is approximately 0.5*(2.7-0)=1.35 volts, which equals VTT, such that the current flow across the 47-ohm resistor of FIG. 17B is approximately zero.

The voltage at the second DQS pin 422 in FIG. 17B is given by V_{DQS2}=2.7*R₄/(R₁+R₂+R₃+R₄)=0.59 volts and the current flowing through the second DQS pin 422 is given by I_{DQS2}=0.59/R₄=34 millamps. The power dissipation in the DQS driver of the second memory device 420 is thus P_{DQS2}=34 mA*0.59 V=20 milliwatts. In contrast, without the first resistor 430 and the second resistor 440, only the 17-ohm impedances of the two memory devices 410, 420 would limit the current flow between the two DQS pins 412, 422, and the power dissipation in the DQS driver of the second memory device 420 would be approximately 107 milliwatts. Therefore, the first resistor 430 and the second resistor 440 of FIGS. 17A and 17B advantageously limit the current flowing between the two memory devices during the time that the DQS pin of one memory device is driven high and the DQS pin of the other memory device is driven low.

In certain embodiments in which there is overshoot or undershoot of the voltages, the amount of current flow can be higher than those expected for nominal voltage values. Therefore, in certain embodiments, the resistances of the first resistor 430 and the second resistor 440 are advantageously selected to account for such overshoot/undershoot of voltages.

For certain such embodiments in which the voltage at the second DQS pin 422 is V_{DQS2}=0.59 volts and the duration of the overdrive condition is approximately 0.8 nanoseconds at maximum, the total surge is approximately 0.59 V*1.2 ns=0.3 V·ns. For comparison, the JEDEC standard for overshoot/undershoot is 2.4 V·ns, so certain embodiments described herein advantageously keep the total surge within predetermined standards (e.g., JEDEC standards).

FIG. 18 schematically illustrates another exemplary memory module 600 compatible with certain embodiments described herein. The memory module 600 comprises a termination bus 605. The memory module 600 further comprises a first memory device 610 having a first data strobe pin 612, a first termination signal pin 614 electrically coupled to the termination bus 605, a first termination circuit 616, and at least one data pin 618. The first termination circuit 616 selectively electrically terminating the first data strobe pin 612 and the first data pin 618 in response to a first signal received by the first termination signal pin 614 from the termination bus 605. The memory module 600 further comprises a second memory device 620 having a second data strobe pin 622 electrically coupled to the first data strobe pin 612, a second termination signal pin 624, a second termination circuit 626, and at least one data pin 628. The second termination signal pin 624 is electrically coupled to a voltage, wherein the second termination circuit 626 is responsive to the voltage by not terminating the second data strobe pin 622 or the second data

US 7,881,150 B2

39

pin 628. The memory module 600 further comprises at least one termination assembly 630 having a third termination signal pin 634, a third termination circuit 636, and at least one termination pin 638 electrically coupled to the data pin 628 of the second memory device 620. The third termination signal pin 634 is electrically coupled to the termination bus 605. The third termination circuit 636 selectively electrically terminates the data pin 628 of the second memory device 620 through the termination pin 638 in response to a second signal received by the third termination signal pin 634 from the termination bus 605.

FIG. 19 schematically illustrates a particular embodiment of the memory module 600 schematically illustrated by FIG. 18. The memory module 600 comprises an on-die termination (ODT) bus 605. The memory module 600 comprises a first memory device 610 having a first data strobe (DQS) pin 612, a first ODT signal pin 614 electrically coupled to the ODT bus 605, a first ODT circuit 616, and at least one data (DQ) pin 618. The first ODT circuit 616 selectively electrically terminates the first DQS pin 612 and the DQ pin 618 of the first memory device 610 in response to an ODT signal received by the first ODT signal pin 614 from the ODT bus 605. This behavior of the first ODT circuit 616 is schematically illustrated in FIG. 14 by the switches 672, 676 which are selectively closed (dash-dot line) or opened (solid line).

The memory module 600 further comprises a second memory device 620 having a second DQS pin 622 electrically coupled to the first DQS pin 612, a second ODT signal pin 624, a second ODT circuit 626, and at least one DQ pin 628. The first DQS pin 612 and the second DQS pin 622 are concurrently active when the first memory device 610 and the second memory device 620 are concurrently enabled. The second ODT signal pin 624 is electrically coupled to a voltage (e.g., ground), wherein the second ODT circuit 626 is responsive to the voltage by not terminating the second DQS pin 622 or the second DQ pin 624. This behavior of the second ODT circuit 626 is schematically illustrated in FIG. 14 by the switches 674, 678 which are opened.

The memory module 600 further comprises at least one termination assembly 630 having a third ODT signal pin 634 electrically coupled to the ODT bus 605, a third ODT circuit 636, and at least one termination pin 638 electrically coupled to the DQ pin 628 of the second memory device 620. The third ODT circuit 636 selectively electrically terminates the DQ pin 628 of the second memory device 620 through the termination pin 638 in response to an ODT signal received by the third ODT signal pin 634 from the ODT bus 605. This behavior of the third ODT circuit 636 is schematically illustrated in FIG. 19 by the switch 680 which is either closed (dash-dot line) or opened (solid line).

In certain embodiments, the termination assembly 630 comprises discrete electrical components which are surface-mounted or embedded on the printed-circuit board of the memory module 600. In certain other embodiments, the termination assembly 630 comprises an integrated circuit mounted on the printed-circuit board of the memory module 600. Persons skilled in the art can provide a termination assembly 630 in accordance with embodiments described herein.

Certain embodiments of the memory module 600 schematically illustrated by FIG. 19 advantageously avoid the problem schematically illustrated by FIG. 12 of electrically connecting the internal termination resistances of the DQS pins of the two memory devices in parallel. As described above in relation to FIG. 14, FIGS. 18 and 19 only show one DQ pin for each memory device for simplicity. Other embodiments have a plurality of DQ pins for each memory device. In

40

certain embodiments, each of the first ODT circuit 616, the second ODT circuit 626, and the third ODT circuit 636 are responsive to a high voltage or signal level by enabling the corresponding termination resistors and are responsive to a low voltage or signal level (e.g., ground) by disabling the corresponding termination resistors. In other embodiments, each of the first ODT circuit 616, the second ODT circuit 626, and the third ODT circuit 636 are responsive to a high voltage or signal level by disabling the corresponding termination resistors and are responsive to a low voltage or signal level (e.g., ground) by enabling the corresponding termination resistors. Furthermore, the switches 672, 674, 676, 678, 680 of FIG. 18 are schematic representations of the enabling and disabling operation of the ODT circuits 616, 626, 636 and do not signify that the ODT circuits 616, 626, 636 necessarily include mechanical switches.

The first ODT signal pin 614 of the first memory device 610 receives an ODT signal from the ODT bus 605. In response to this ODT signal, the first ODT circuit 616 selectively enables or disables the termination resistance for both the first DQS pin 612 and the DQ pin 618 of the first memory device 610. The second ODT signal pin 624 of the second memory device 620 is tied (e.g., directly hard-wired) to the voltage (e.g., ground), thereby disabling the internal termination resistors 654, 658 on the second DQS pin 622 and the second DQ pin 628, respectively, of the second memory device 620 (schematically shown by open switches 674, 678 in FIG. 19). The second DQS pin 622 is electrically coupled to the first DQS pin 612, so the termination resistance for both the first DQS pin 612 and the second DQS pin 622 is provided by the termination resistor 652 internal to the first memory device 510.

The termination resistor 656 of the DQ pin 618 of the first memory device 610 is enabled or disabled by the ODT signal received by the first ODT signal pin 614 of the first memory device 610 from the ODT bus 605. The termination resistance of the DQ pin 628 of the second memory device 620 is enabled or disabled by the ODT signal received by the third ODT signal pin 634 of the termination assembly 630 which is external to the second memory device 620. Thus, in certain embodiments, the first ODT signal pin 614 and the third ODT signal pin 634 receive the same ODT signal from the ODT bus 605, and the termination resistances for both the first memory device 610 and the second memory device 620 are selectively enabled or disabled in response thereto when these memory devices are concurrently enabled. In this way, certain embodiments of the memory module 600 schematically illustrated by FIG. 19 provides external or off-chip termination of the second memory device 620.

Certain embodiments of the memory module 600 schematically illustrated by FIG. 19 advantageously allow the use of two lower-cost readily-available 512-Mb DDR-2 SDRAM devices to provide the capabilities of a more expensive 1-GB DDR-2 SDRAM device. Certain such embodiments advantageously reduce the total cost of the resultant memory module 600.

Certain embodiments described herein advantageously increase the memory capacity or memory density per memory slot or socket on the system board of the computer system. Certain embodiments advantageously allow for higher memory capacity in systems with limited memory slots. Certain embodiments advantageously allow for flexibility in system board design by allowing the memory module 10 to be used with computer systems designed for different numbers of ranks (e.g., either with computer systems designed for two-rank memory modules or with computer

US 7,881,150 B2

41

systems designed for four-rank memory modules). Certain embodiments advantageously provide lower costs of board designs.

In certain embodiments, the memory density of a memory module is advantageously doubled by providing twice as many memory devices as would otherwise be provided. For example, pairs of lower-density memory devices can be substituted for individual higher-density memory devices to reduce costs or to increase performance. As another example, twice the number of memory devices can be used to produce a higher-density memory configuration of the memory module. Each of these examples can be limited by the number of chip select signals which are available from the memory controller or by the size of the memory devices. Certain embodiments described herein advantageously provide a logic mechanism to overcome such limitations.

Various embodiments of the present invention have been described above. Although this invention has been described with reference to these specific embodiments, the descriptions are intended to be illustrative of the invention and are not intended to be limiting. Various modifications and applications may occur to those skilled in the art without departing from the true spirit and scope of the invention.

What is claimed is:

1. A circuit configured to be mounted on a memory module so as to be electrically coupled to a plurality of double-data-rate (DDR) memory devices arranged in one or more ranks on the memory module, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals, comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;
a register;

a phase-lock loop device configured to be operationally coupled to the plurality of DDR memory devices, the logic element, and the register,

wherein the circuit is configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the plurality of DDR memory devices, wherein the system memory domain has a first memory density per rank and the physical memory domain has a second memory density per rank less than the first memory density per rank.

2. The circuit of claim 1, wherein the plurality of DDR memory devices has a first number of DDR memory devices, and the circuit is configurable to selectively isolate the loads of a second number of the DDR memory devices from the computer system, the second number less than the first number.

3. The circuit of claim 2, wherein the circuit is configurable to selectively isolate the loads of the second number of the DDR memory devices from the computer system in response to a command or address signal of the set of input signals.

4. The circuit of claim 1, wherein the system memory domain is compatible with a first number of chip-select signals, and the physical memory domain is compatible with a second number of chip-select signals greater than the first number of chip-select signals.

5. The circuit of claim 1, wherein two or more of the phase-lock loop device, the logic element, and the register are portions of a single component.

42

6. The circuit of claim 1, wherein the circuit is configurable to transmit a set of output signals to the plurality of DDR memory devices, the set of output signals compatible with the physical memory domain.

5 7. The circuit of claim 6, wherein the circuit is configurable to receive a command signal from the computer system, select at least one rank of the one or more ranks, and transmit the command signal to at least one DDR memory device of the selected rank of DDR memory devices.

10 8. The circuit of claim 6, wherein the set of input signals comprises a first set of address and command signals and the set of output signals comprises a second set of address and command signals.

15 9. The circuit of claim 8, wherein the circuit is configurable to store an address signal of the set of input signals during a row access procedure and pass the stored address signal as an address signal of the set of output signals during a subsequent column address procedure.

10 10. The circuit of claim 1, further configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.

15 11. The circuit of claim 1, further configurable to receive an address signal of the set of input signals and to use the address signal to translate between the system memory domain and the physical memory domain.

12. The circuit of claim 11, wherein the address signal is a row address bit and the circuit is configurable to latch the row address bit.

30 13. The circuit of claim 11, wherein the address signal is a row address bit and the circuit is configurable to use the row address bit and a first number of chip-select signals of the set of input signals to generate a second number of chip-select signals greater than the first number of chip-select signals.

35 14. The circuit of claim 11, wherein the address signal is a row address bit and the circuit is configurable to use the row address bit and a first number of column address strobe (CAS) signals to generate a second number of CAS signals greater than the first number of CAS signals.

40 15. A circuit configured to be mounted on a memory module so as to be electrically coupled to a first double-data-rate (DDR) memory device having a first data signal line and a first data strobe line, to a second DDR memory device having a second data signal line and a second data strobe line, and to a common data signal line, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;
a register;

a phase-lock loop device configured to be operationally coupled to the first DDR memory device, the second DDR memory device, the logic element, and the register, wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the memory module, wherein the system memory domain has a first memory density per memory device, and the physical memory domain has a second

US 7,881,150 B2

43

memory density per memory device less than the first memory density per memory device.

16. The circuit of claim 15, wherein the two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.

17. The circuit of claim 16, further comprising one or more switches selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the one or more switches operatively coupled to the logic element to receive control signals from the logic element.

18. The circuit of claim 15, wherein the memory module further comprises a third data strobe signal line connectable to the computer system, the circuit electrically coupled to the first data strobe signal line, to the second data strobe signal line, and to the third data strobe signal line, the circuit configurable to selectively electrically couple the first data strobe signal line to the third data strobe signal line and to selectively electrically couple the second data strobe signal line to the third data strobe signal line.

19. The circuit of claim 15, wherein the circuit is configurable to save an address signal of the set of input signals during a row access procedure and pass the saved address signal as an address signal during a subsequent column address procedure.

20. The circuit of claim 15, further configurable to be responsive to the set of input signals by selectively electrically coupling both the first data signal line and the second data signal line to the computer system.

21. The circuit of claim 15, wherein the circuit is configurable to selectively isolate both the first data signal line and the second data signal line from the computer system.

22. A circuit configured to be mounted on a memory module so as to be electrically coupled to a plurality of double-data-rate (DDR) memory devices arranged in one or more ranks on the memory module, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals, comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;

a register;

a phase-lock loop device configured to be operationally coupled to the plurality of DDR memory devices, the logic element, and the register,

wherein the circuit is configurable to be responsive to the set of input signals by selectively isolating one or more loads of the DDR memory devices from the computer system, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the plurality of DDR memory devices, wherein the system memory domain is compatible with a first number of chip-select signals, and the physical memory domain is compatible with a second number of chip-select signals greater than the first number of chip-select signals.

23. The circuit of claim 22, wherein the plurality of DDR memory devices has a first number of DDR memory devices, and the circuit is configurable to selectively isolate the loads of a second number of the DDR memory devices from the computer system, the second number less than the first number.

44

24. The circuit of claim 22, wherein two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.

25. The circuit of claim 22, wherein the circuit is configurable to store an address signal of the set of input signals during a row access procedure and to pass the stored address signal during a subsequent column address procedure.

26. The circuit of claim 22, further configurable to selectively isolate a data signal line of a DDR memory device of the plurality of DDR memory devices from the computer system.

27. The circuit of claim 22, further configurable to receive an address signal of the set of input signals and to use the address signal to translate between the system memory domain and the physical memory domain.

28. The circuit of claim 27, wherein the address signal is a row address bit and the circuit is configurable to latch the row address bit.

29. The circuit of claim 27, wherein the address signal is a row address bit and the circuit is configurable to use the row address bit and the first number of chip-select signals to generate the second number of chip-select signals.

30. The circuit of claim 27, wherein the address signal is a row address bit and the circuit is configurable to use the row address bit and a first number of column address strobe (CAS) signals to generate a second number of CAS signals greater than the first number of CAS signals.

31. A circuit configured to be mounted on a memory module so as to be electrically coupled to a first double-data-rate (DDR) memory device having a first data signal line and a first data strobe line, to a second DDR memory device having a second data signal line and a second data strobe line, and to a common data signal line, the memory module configured to be electrically coupled to a memory controller of a computer system so as to receive a set of input signals comprising row address signals, column address signals, bank address signals, and chip-select signals, the set of input signals compatible with a system memory domain of the computer system, the circuit comprising:

a logic element;

a register;

a phase-lock loop device configured to be operationally coupled to the first DDR memory device, the second DDR memory device, the logic element, and the register, wherein the circuit is configurable to be responsive to the set of input signals by selectively electrically coupling the first data signal line to the common data signal line and selectively electrically coupling the second data signal line to the common data signal line, the circuit configurable to translate between the system memory domain of the computer system and a physical memory domain of the memory module, wherein the system memory domain is compatible with a first number of chip-select signals, and the physical memory domain is compatible with a second number of chip-select signals greater than the first number of chip-select signals.

32. The circuit of claim 31, wherein two or more of the logic element, the register, and the phase-lock loop device are portions of a single component.

33. The circuit of claim 31, further comprising one or more switches configurable to selectively electrically couple the first data signal line to the common data signal line and to selectively electrically couple the second data signal line to the common data signal line, the one or more switches operatively coupled to the logic element to receive control signals from the logic element.

US 7,881,150 B2

45

34. The circuit of claim 31, wherein the circuit is configurable to save an address signal of the set of input signals during a row access procedure and to pass the saved address signal as an address signal during a subsequent column address procedure.

35. The circuit of claim 31, further configurable to be responsive to the set of input signals by selectively electri-

46

cally coupling both the first data signal line and the second data signal line to the computer system.

36. The circuit of claim 31, further configurable to selectively isolate both the first data signal line and the second data signal line from the computer system.

* * * * *



US008081536B1

(12) **United States Patent**
Solomon et al.

(10) **Patent No.:** US 8,081,536 B1
(45) **Date of Patent:** *Dec. 20, 2011

(54) **CIRCUIT FOR MEMORY MODULE**

(75) Inventors: **Jeffrey C. Solomon**, Irvine, CA (US);
Jayesh R. Bhakta, Cerritos, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **13/032,470**

(22) Filed: **Feb. 22, 2011**

Related U.S. Application Data

(63) Continuation of application No. 12/955,711, filed on Nov. 29, 2010, now Pat. No. 7,916,574, which is a continuation of application No. 12/629,827, filed on Dec. 2, 2009, now Pat. No. 7,881,150, which is a continuation of application No. 12/408,652, filed on Mar. 20, 2009, now Pat. No. 7,636,274, which is a continuation of application No. 11/335,875, filed on Jan. 19, 2006, now Pat. No. 7,532,537, which is a continuation-in-part of application No. 11/173,175, filed on Jul. 1, 2005, now Pat. No. 7,289,386, which is a continuation-in-part of application No. 11/075,395, filed on Mar. 7, 2005, now Pat. No. 7,286,436.

(60) Provisional application No. 60/645,087, filed on Jan. 19, 2005, provisional application No. 60/588,244, filed on Jul. 15, 2004, provisional application No. 60/550,668, filed on Mar. 5, 2004, provisional application No. 60/575,595, filed on May 28, 2004, provisional application No. 60/590,038, filed on Jul. 21, 2004.

(51) **Int. Cl.**

GIIC 8/00 (2006.01)

(52) **U.S. Cl.** **365/230.06; 365/51; 365/149;**
365/230.03; 365/233.13; 711/105; 711/115

(58) **Field of Classification Search** 365/51,
365/149, 230.03, 230.06, 233.13; 711/105,
711/115
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,368,515 A 1/1983 Nielsen
(Continued)

FOREIGN PATENT DOCUMENTS

JP 10320270 A * 12/1998
(Continued)

OTHER PUBLICATIONS

English Abstract for JP 10320270A retrieved from eSpacenet on Oct. 18, 2011.*

(Continued)

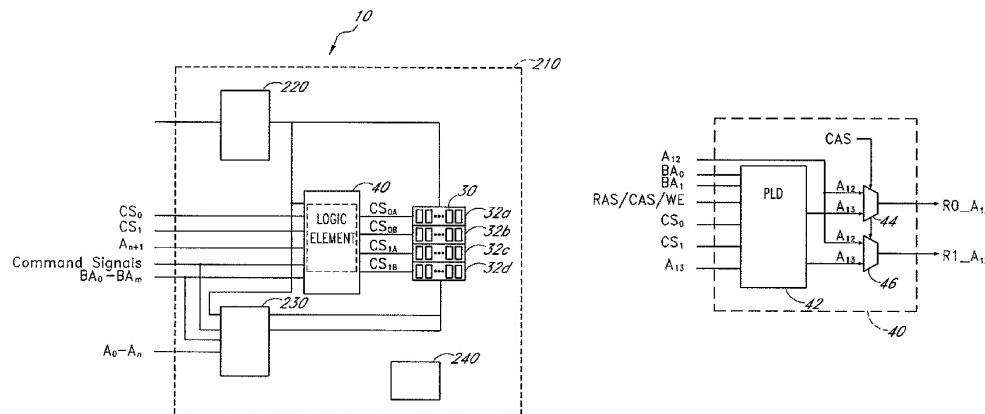
Primary Examiner — Alexander Sofocleous

(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson & Bear LLP

ABSTRACT

A circuit is configured to be mounted on a memory module configured to be operationally coupled to a computer system. The memory module has a first number of ranks of double-data-rate (DDR) memory circuits activated by a first number of chip-select signals. The circuit is configurable to receive a set of signals comprising address signals and a second number of chip-select signals smaller than the first number of chip-select signals. The circuit is further configurable to generate phase-locked clock signals, to selectively isolate a load of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals, and to generate the first number of chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of chip-select signals.

42 Claims, 23 Drawing Sheets



US 8,081,536 B1

Page 2

U.S. PATENT DOCUMENTS	
4,392,212 A	7/1983 Miyasaka et al.
4,633,429 A	12/1986 Lewandowski et al.
4,670,748 A	6/1987 Williams
4,866,603 A	9/1989 Chiba
4,958,322 A	9/1990 Kosugi et al.
4,961,172 A	10/1990 Shubat et al.
4,980,850 A	12/1990 Morgan
5,247,643 A	9/1993 Shottan
5,345,412 A	9/1994 Shiratsuchi
5,426,753 A	6/1995 Moon
5,483,497 A	1/1996 Mochizuki et al.
5,495,435 A	2/1996 Sugahara
5,581,498 A	12/1996 Ludwig et al.
5,590,071 A	12/1996 Kolor et al.
5,699,542 A	12/1997 Mehta et al.
5,702,984 A	12/1997 Bertin et al.
5,703,826 A	12/1997 Hush et al.
5,745,914 A	4/1998 Connolly et al.
5,802,395 A	9/1998 Connolly et al.
5,805,520 A	9/1998 Anglada et al.
5,822,251 A	10/1998 Bruce et al.
RE36,229 E	6/1999 Cady
5,926,827 A	7/1999 Dell et al.
5,959,930 A	9/1999 Sakurai
5,963,464 A	10/1999 Dell et al.
5,966,736 A	10/1999 Gittinger et al.
6,018,787 A	1/2000 Ip
6,044,032 A	3/2000 Li
6,070,217 A	5/2000 Connolly et al.
6,070,227 A	5/2000 Rokicki
6,097,652 A	8/2000 Roh
6,108,745 A	8/2000 Gupta et al.
6,134,638 A	10/2000 Olarig et al.
6,151,271 A	11/2000 Lee
6,154,418 A	11/2000 Li
6,154,419 A	11/2000 Shakkarwar
6,185,654 B1	2/2001 Van Doren
6,205,516 B1	3/2001 Usami
6,209,074 B1	3/2001 Dell et al.
6,226,709 B1	5/2001 Goodwin et al.
6,226,736 B1	5/2001 Niot
6,233,650 B1	5/2001 Johnson et al.
6,247,088 B1	6/2001 Seo et al.
6,317,352 B1	11/2001 Halbert et al.
6,400,637 B1	6/2002 Akamatsu et al.
6,408,356 B1	6/2002 Dell
6,414,868 B1	7/2002 Wong et al.
6,415,374 B1	7/2002 Faue et al.
6,446,158 B1	9/2002 Karabatsos
6,446,184 B2	9/2002 Dell et al.
6,453,381 B1	9/2002 Yuan et al.
6,470,417 B1	10/2002 Kolor et al.
6,480,439 B2	11/2002 Tokutome et al.
6,502,161 B1	12/2002 Perego et al.
6,518,794 B2	2/2003 Coteus et al.
6,526,473 B1	2/2003 Kim
6,530,007 B2	3/2003 Olarig et al.
6,530,033 B1	3/2003 Raynham et al.
6,553,450 B1	4/2003 Dodd et al.
6,618,320 B2	9/2003 Hasegawa et al.
6,621,496 B1	9/2003 Ryan
6,625,081 B2	9/2003 Roohparvar et al.
6,625,687 B1	9/2003 Halbert et al.
6,636,935 B1	10/2003 Ware et al.
6,646,949 B1	11/2003 Ellis et al.
6,658,509 B1	12/2003 Bonella et al.
6,674,684 B1	1/2004 Shen
6,681,301 B1	1/2004 Mehta et al.
6,683,372 B1	1/2004 Wong et al.
6,697,888 B1	2/2004 Halbert et al.
6,705,877 B1	3/2004 Li et al.
6,717,855 B2	4/2004 Lai
6,738,880 B2	5/2004 Lai et al.
6,742,098 B1	5/2004 Halbert et al.
6,754,797 B2	6/2004 Wu et al.
6,785,189 B2	8/2004 Jacobs et al.
6,788,592 B2	9/2004 Nakata et al.
6,807,125 B2	10/2004 Coteus et al.
	6,807,650 B2 10/2004 Lamb et al.
	6,813,196 B2 11/2004 Park et al.
	6,834,014 B2 12/2004 Yoo et al.
	6,854,042 B1 2/2005 Karabatsos
	6,880,094 B2 4/2005 LaBerge
	6,889,304 B2 5/2005 Perego et al.
	6,912,615 B2 6/2005 Nicolai
	6,912,628 B2 6/2005 Wicki et al.
	6,925,028 B2 8/2005 Hosokawa et al.
	6,944,694 B2 9/2005 Pax
	6,950,366 B1 9/2005 Lapidus et al.
	6,961,281 B2 11/2005 Wong et al.
	6,981,089 B2 12/2005 Dodd et al.
	6,982,892 B2 1/2006 Lee et al.
	6,982,893 B2 1/2006 Jakobs
	6,990,043 B2 1/2006 Kuroda et al.
	6,996,686 B2 2/2006 Doblar et al.
	7,007,130 B1 2/2006 Holman
	7,007,175 B2 2/2006 Chang et al.
	7,046,538 B2 5/2006 Kinsley et al.
	7,054,179 B2 5/2006 Cogdill et al.
	7,065,626 B2 6/2006 Schumacher et al.
	7,073,041 B2 7/2006 Dwyer et al.
	7,078,793 B2 7/2006 Ruckerbauer et al.
	7,120,727 B2 10/2006 Lee et al.
	7,124,260 B2 10/2006 LaBerge et al.
	7,127,584 B1 10/2006 Thompson et al.
	7,130,952 B2 10/2006 Nanki et al.
	7,133,960 B1 11/2006 Thompson et al.
	7,133,972 B2 11/2006 Jeddelloh
	7,142,461 B2 11/2006 Janzen
	7,149,841 B2 12/2006 LaBerge
	7,167,967 B2 1/2007 Bungo et al.
	7,181,591 B2 2/2007 Tsai
	7,200,021 B2 4/2007 Raghuram
	7,227,910 B2 6/2007 Lipka
	7,266,639 B2 9/2007 Raghuram
	7,272,709 B2 9/2007 Zitlaw et al.
	7,281,079 B2 10/2007 Bains et al.
	7,286,436 B2 10/2007 Bhakta et al.
	7,289,386 B2 10/2007 Bhakta et al.
	7,346,750 B2 3/2008 Ishikawa
	7,356,639 B2 4/2008 Perego et al.
	7,370,238 B2 5/2008 Billlick et al.
	7,437,591 B1 10/2008 Wong
	7,461,182 B2 12/2008 Fukushima et al.
	7,471,538 B2 12/2008 Hofstra
	7,532,537 B2 5/2009 Solomon et al.
	7,619,912 B2 11/2009 Bhakta et al.
	7,636,274 B2 12/2009 Solomon et al.
	7,864,627 B2 1/2011 Bhakta et al.
	7,881,150 B2 2/2011 Solomon et al.
	7,916,574 B1 3/2011 Solomon et al.
	8,001,434 B1 * 8/2011 Lee et al. 714/733
	2001/0003198 A1 6/2001 Wu
	2001/0052057 A1 12/2001 Lai
	2002/008633 A1 7/2002 Kong et al.
	2003/0063514 A1 4/2003 Faue
	2003/0090359 A1 5/2003 OK
	2003/0090879 A1 5/2003 Doblar et al.
	2003/0191995 A1 10/2003 Abrosimov et al.
	2003/0210575 A1 11/2003 Seo et al.
	2004/0000708 A1 1/2004 Rapport et al.
	2004/0037158 A1 2/2004 Coteus et al.
	2004/0201968 A1 10/2004 Tafolla
	2005/0036378 A1 2/2005 Kawaguchi et al.
	2005/0281096 A1 12/2005 Bhakta et al.
	2006/0044860 A1 3/2006 Kinsley et al.
	2006/0117152 A1 6/2006 Amidi
	2006/0126369 A1 6/2006 Raghuram
	2006/0129755 A1 6/2006 Raghuram
	2006/0179206 A1 8/2006 Brittain et al.
	2006/0259711 A1 11/2006 Oh
	2006/0267172 A1 11/2006 Nguyen et al.
	2006/0277355 A1 12/2006 Ellsberry et al.
	2010/0128507 A1 5/2010 Solomon et al.
	2011/0016250 A1 1/2011 Lee et al.

US 8,081,536 B1

Page 3

2011/0016269 A1	1/2011	Lee et al.
2011/0090749 A1*	4/2011	Bhakta et al. 365/191
2011/0125966 A1	5/2011	Amidi et al.

FOREIGN PATENT DOCUMENTS

WO	WO 92/02879	2/1992
WO	WO 94/07242	3/1994
WO	WO 95/34030	12/1995
WO	WO 02/058069	7/2002
WO	WO 03/017283	2/2003
WO	WO 03/069484	8/2003
WO	WO 2006/055497	5/2006

OTHER PUBLICATIONS

English Machine Translation of JP 10320270A retrieved from PAJ on Oct. 18, 2011.*
U.S. Appl. No. 12/815,339, filed Jun. 14, 2010; Owned by Netlist, Inc.
U.S. Appl. No. 12/774,632, filed May 5, 2010; Owned by Netlist, Inc.
U.S. Appl. No. 12/422,853, filed Apr. 13, 2009; Owned by Netlist, Inc.
U.S. Appl. No. 12/981,380, filed Nov. 29, 2010; Owned by Netlist, Inc.
U.S. Appl. No. 12/912,623, filed Oct. 26, 2010; Owned by Netlist, Inc.
Reexam U.S. Appl. No. 95/001,339, filed Jun. 8, 2010; Owned by Netlist, Inc.
Reexam U.S. Appl. No. 95/001,337, filed Jun. 4, 2010; Owned by Netlist, Inc.
Reexam U.S. Appl. No. 95/000,546, filed May 11, 2010; Owned by Netlist, Inc.
Reexam U.S. Appl. No. 95/001,381, filed Jun. 9, 2010; Owned by Netlist, Inc.
“64 & 72 Pin Zip/Simm Sram Module”, JEDEC, Standard No. 21-C, www.jedec.com/download/search/4_04_01.pdf, Jun. 1997 pp. 4.4-1-1.
“Bank Striping of Data Across Internal SDRAM Banks,” IP.com, IPCOM000013697D, 2000.
“Distributed Memory Mapping,” IP.com, IPCOM000014788D, 2000.
“Information Huawei or FPGA-Take Five,” Electronic News, 2002, p. 24.
“Method for a high-performance DRAM address mapping mechanism,” IP.com, IPCOM00008164D, 2002.
“Method for memory probing on a multiple-DIMM bus,” IP.com, IPCOM000019063D, 2003.
Reexam U.S. Appl. No. 95/000,577 for U.S. Pat. No. 7,289,386, filed Oct. 20, 2010.
Reexam U.S. Appl. No. 95/000,578 for U.S. Pat. No. 7,619,912, filed Oct. 20, 2010.
Reexam U.S. Appl. No. 95/000,579 for U.S. Pat. No. 7,619,912, filed Oct. 21, 2010.
“Method for multiple device interface testing using a single device,” IP.com, IPCOM000010054D, 2002.
“Quad Band Memory (QBMA™): DDR200/266/333 devices producing DDR400/533/667” (the “QBMA Reference”), published by the QBMA Alliance, Platform Conference, San Jose, California, Jan. 23-24, 2002.
“DDR SDRAM RDIMM Features,” Micron Technology, Inc., 2002.
“PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification” JEDEC, Standard No. 21-C, Revision 1-3, Jan. 2002, pp. 4.20.4-1.
“Quad Band Memory (QBMTM): DDR 200/266/333 devices producing DDR 400/533/667,” Platform Conference, Jan. 23-24, 2002.
Abali, B. “Memory Expansion Technology (MXT): Software Support and Performance,” IBM J. Res. & Dev., vol. 45, No. 2, 2001, pp. 287-300.
Arlington, DL Evans. “Enhancement of Memory Card Redundant Bit Usage Via Simplified Fault Alignment Exclusion,” IBM Technical Disclosure Bulletin, 1987.
Arroyo et al. “Method of executing Manufacturing ROM Code Without Removing System Roms,” IP.com, IPCOM000037214D, 1989.
Barr, Michael. “Programmable Logic: What’s it to Ya?,” Embedded Systems Programming, Jun. 1999, pp. 75-84.
Bennayoun et al. “Input/Output Chip Select Doubler,” IBM Technical Disclosure Bulletin, vol. 38, No. 04 1995, pp. 237-240.
Blum et al. “Fast Multichip Memory System With Power Select Signal,” IBM Technical Disclosure Bulletin, 1979.
Carvalho, Carlos; “The Gap between Processor and Memory Speeds”; ICCA’02, published 2002.
Cuppu et al. “Concurrency, Latency, or System Overhead: Which Has the Largest Impact on Uniprocessor DRAM-System Performance?,” IEEE, 2001, pp. 62-71.
Cuppu et al. “High-Performance DRAMs in Workstation Environments,” IEEE Transactions on Computers, vol. 50, No. 11, 2001, pp. 1133-1153.
Cuppu et al. “A Performance Comparison of Contemporary DRAM Architectures,” *IEEE Proceedings of the 26th International Symposium on Computer Architectures*, May 2-4, 1999, Atlanta, Georgia, pp. 1-12.
Denneau, M. “Logic Processor for Logic Simulation Machine,” IBM Technical Disclosure Bulletin, vol. 25, No. 1, 1982.
Fairchild Semiconductor. “DM74LS138 DM74LS139 Decoder/Demultiplexer,” Fairchild Semiconductor Corporation, 2000.
Fitzgerald et al. “Chip Select Circuit for Multi-Chip RAM Modules,” IP.com, IPCOM00004404D, 1984.
Freedman, Alan. “The Computer Glossary,” The Complete Illustrated Dictionary, American Management Association, 2001.
Google, Inc. v. Netlist, Inc., No. 4:08-cv-04144-SBA, Netlist Inc.’s Answer to Complaint and Counterclaim (N. D. Ca. Filed Nov. 18, 2008).
Google, Inc. v. Netlist, Inc., No. C 08-04144 SBA Google Inc.’s Invalidity Contentions Pursuant to PAT. L.F. 3-3, dated Apr. 13, 2009.
Google, Inc. v. Netlist, Inc., No. C08 04144, Complaint for Declaratory Relief, (N. D. Ca. Dated Aug. 29, 2008).
Gray, KS. “Fet Ram Chip Double Density Scheme,” IP.com, IPCOM000043942D, 1984.
Grimes et al. “Access Rate/Availability Improvement Logic for Dynamic Memories,” IBM Technical Disclosure Bulletin, Oct. 1982.
Gupta et al. “Designing and Implementing a Fast Crossbar Scheduler,” IEEE Micro, 1999, pp. 20-28.
Hession et al. “Chip Select Technique for Multi Chip Decoding,” IP.com, IPCOM000070404D, 1985.
Hewlett-Packard. “Memory technology evolution: an overview of system memory technologies,” technology brief, 7th edition. 2003.
Hoare et al. “An 88-Way Multiprocessor Within an FPGA With Customizable Instructions,” Proceedings of the 18th International Parallel and Distributed Processing Symposium, 2004.
Intel Corporation, 66/100 MHz PC SDRAM 64-Bit Non-ECC/Parity 144 Pin Unbuffered SO-DIMM Specification, Revision 1.0, Feb. 1999.
Intel Corporation, PC SDRAM Registered DIMM Design Support Document, Revision 1.2, Oct. 1998.
Jacob, Bruce L.; “Synchronous DRAM Architectures, Organizations, and Alternative Technologies”. University of Maryland, Dec. 10, 2002.
JEDEC “JEDEC Standard: Double Data Rate (DDR) SDRAM Specification”, JESD79C Mar. 2003.
JEDEC Standard JESD79D, “Double Data Rate (DDR) SDRAM Specification,” published Feb. 2004.
JEDEC Standard No. 21-C, “PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification,” Revision 1.3, Jan. 2002.
JEDEC Standard No. 21-C, 4.20.5-184 Pin. PC1600/2100 DDR SDRAM Unbuffered DIMM Design Specification, Revision 1.1, Release 11 b. Published Apr. 2003.
JEDEC Standard No. 21-C, 4.20.5-184 Pin. PC2700/PC2100/PC1600 DDR SDRAM Unbuffered SO-DIMM Reference Design Specification, Revision 1.1, Release 11b, Apr. 26, 2002.
JEDEC Standard No. 21-C, 4.20.2—168 Pin, PC133 SDRAM Registered Design Specification, Revision 1.4, Release 11a, Feb. 2002.
JEDEC Standard No. 21-C, 4.20.3—144 Pin, PC133 SDRAM Unbuffered SO-DIMM, Reference Design Specification, Revision 1.0, Release 11. Published Oct. 2003.

US 8,081,536 B1

Page 4

- JEDEC Standard No. 21-C, DDR SDRAM PC2100 and PC1600 DDR SDRAM Registered DIMM Design Specification, Revision 1.3, Release 11 b, Jan. 2002.
- JEDEC Standard, "Definition of the SSTV16859 2.5 V 13-Bit to 26-Bit SSTL_2 Registered Buffer for Stacked DDR DIMM Applications," JESD82-4B, May 2003.
- Jin et al. "Embedded Memory in System-On-Chip Design: Architecture and Prototype Implementation," CCECE, 2003, pp. 141-146.
- Jin et al. "Prototype Implementation and Evaluation of a Multibank Embedded Memory Architecture in Programmable Logic," IEEE, 2003, pp. 13-16.
- Kane et al. "Read Only Store Memory Extension," IP.com, IPCOM000082845D, 1975.
- Karabatsos, C., "Quad Band Memory (QBM) Technology", Kentron Technologies, Inc., Apr. 2001, pp. 1-5.
- Kellogg, Mark; "PC133: SDRAM Main Memory Performance Reaches New Heights"; IBM Microelectronics, 1999.
- Keltcher et al.; "The AMD Opteron Processor for Multiprocessor Servers"; IEEE Micro, vol. 23, No. 2, Mar. 2003.
- Kirihata et al.; "A 390-mm, 16-Bank, 1-Gb DDR SDRAM with Hybrid Bitline Architecture"; IEEE Journal of Solid-State Circuits, vol. 34, No. 11, Nov. 1999.
- Kornaros et al. "A Fully-Programmable Memory Management System Optimizing Queue Handling at Multi Gigabit Rates," DAC, 2003, pp. 54-59.
- Lee et al. "A banked-promotion translation lookaside buffer system," Journal of Systems Architecture, vol. 47, 2002, pp. 1065-1078.
- Lee et al. "An on-chip cache compression technique to reduce decompression overhead and design complexity." Journal of Systems Architecture, vol. 46, 2000, pp. 1365-1382.
- Letter from G. Hopkins Guy III, Orrick, Herrington & Sutcliffe LLP, to R. Scott Oliver, Morrison & Foerster, (Apr. 14, 2009).
- Lin et al. "Designing a Modern Memory Hierarchy with Hardware Prefetching," IEEE Transactions on Computers, vol. 50, No. 11, 2001, pp. 1202-1217.
- Luthra et al. "Interface Synthesis Using Memory Mapping for an FPGA Platform," Proceedings of the 21st International Conference on Computer Design, 2003.
- Matick et al. "Read-Select Capability for Static Random-Access Memory," IBM Technical Disclosure Bulletin, 1985, pp. 6640-6642.
- Matick, RE. "Logic and Decoder Arrangement for Controlling Spill/Wrap Boundaries of a Bit-Addressable Memory Decoder," IBM Technical Disclosure Bulletin, 1984.
- MetaRAM, Inc. v. Netlist, Inc.* No. 3:09-cv-01309-VRW, MetaRAM's Reply to Netlist's Counterclaims, (N.D. Ca. Filed Jun. 3, 2009).
- MetaRAM, Inc. v. Netlist, Inc.*, No. 3:09-cv-01309-VRW, Netlist's Answer to Complaint and Counterclaims, (N.D. Ca, filed May 11, 2009).
- MetaRAM, Inc. v. Netlist, Inc.*, No. C09 01309, Complaint for Patent Infringement, (N.D. Ca. Filed Mar. 25, 2009).
- Meyers et al. "Use of Partially Good Memory Chips," IP.com, IPCOM0000066246D, 1979.
- Miles J. Murdocca et al., "Principles of Computer Architecture", Prentice Hall, 2000, pp. 249-251.
- Netlist, Inc. v MetaRAM, Inc.*, No. 09-165-GMS, MetaRAM, Inc.'s Answer and Affirmative Defenses to Plaintiffs Complaint, dated Apr. 20, 2009.
- Netlist, Inc. v. MetaRAM, Inc.*, No. 1:09-ccv-00165-GMS, Complaint for Patent Infringement, (D. Del. Filed Mar. 12, 2009).
- Ofek et al. "Partial Two Way Mapping Technique," IBM Technical Disclosure Bulletin, 1969.
- Paldan, David. "Programmable Memory Address Decoding for Microprocessor Memory Device," IP.com, IPCOM000005486D, 1983.
- Pellinger et al. "Dual Addressable Memory," IP.com, IPCOM0000068610D, 1978.
- Plotnick et al. "Shuffle Your Chips for Better Performance," PC Week, 1998, p. 90.
- Schubert et al. "Accelerating system integration by enhancing hardware, firmware, and co-simulation," IBM J. Res. & Dev, vol. 48, No. 3/4, May/Jul. 2004, pp. 569-581.
- Skelton, MH. "Program Controlled Paging Scheme for Memory Expansion," IP.com. IPCOM000050954D, 1982.
- Slegel et al. "IBM's S/390 G5 Microprocessor Design," IEEE Micro, 1999, pp. 12-23.
- Smith, BA. "Chip Select Decoder Circuit," IP.com, IPCOM000063400D, 1985.
- Stelzer, KC. "Planar Memory Boundary Registers with Remap Feature, IBM Technical Disclosure Bulletin," 1993.
- Sunaga et al. "An Enable Signal Circuit for Multiple Small Banks," IP.com, IPCOM000015887D, 2002.
- Sunaga et al. "Continuou RAS Access Method in Multiple-bank DRAM Chip," IP.com, IPCOM000123375D, 1998.
- Toal et al. "A 32-Bit SoPC Implementation of a P5." Proceedings of the Eighth IEEE International Symposium on Computers and Communications, 2003, pp. 1530-1546.
- Tudruj, Marek. "Dynamically reconfigurable heterogenous multiprocessor systems with transputer-controlled communication," Journal of Systems Architecture, vol. 43, 1997, pp. 27-32.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Complaint for Patent Infringement, filed Sep. 22, 2009 in 10 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Defendant Inphi Corporation's Answer to Plaintiffs Complaint for Patent Infringement, filed Nov. 12, 2009 in 6 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Defendant Inphi Corporation's Answer to Plaintiffs First Amended Complaint for Patent Infringement, filed Feb. 11, 2010 in 9 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Defendant Inphi Corporation's Notice of Motion and Motion for Stay Pending Reexaminations and Interference Proceeding Regarding The Patents-In-Suit; Memorandum of Points and Authorities in Support Thereof, filed Apr. 21, 2010 in 28 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Plaintiff Netlist Inc's Opposition to Defendant Inphi Corporation's Motion for Stay Pending Reexaminations and Interference Proceedings Regarding the Patents-In-Suit, filed May 3, 2010 in 23 pages.
- U.S. District Court Central District of California, Case No. CV09 06900, *Netlist, Inc. vs. Inphi Corporation*, Plaintiff Netlist, Inc's First Amended Complaint for Patent Infringement, filed Dec. 23, 2009 in 8 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, [Redacted] Google Inc.'s Responsive Claim Construction Brief, filed Aug. 25, 2009 in 30 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Amended Exhibit A to Joint Claim Construction and Prehearing Statement, filed Oct. 28, 2009 in 1 page.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Appendix 1 to Google's Responsive Claim Construction Brief, filed Aug. 25, 2009 in 4 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Attachment 1 to Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 7 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Attachment 2 to Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 12 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Complaint for Declaratory Relief, filed Aug. 29, 2008 in 49 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Defendant Netlist, Inc's Claim Construction Reply Brief, filed Sep. 22, 2009 in 19 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Defendant Netlist, Inc.'s Opening Claim Construction Brief, filed Jul. 29, 2009 in 21 pages.
- U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Defendant Netlist, Inc.'s Opposition to Google Inc's Motion for Summary Judgment of Invalidity, filed Jul. 6, 2010 in 13 pages.

US 8,081,536 B1

Page 5

U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Exhibit A to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 2 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Exhibit B to Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 36 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Joint Claim Construction and Prehearing Statement, filed Jun. 12, 2009 in 5 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Netlist, Inc.'s Answer to Complaint and Counterclaims, filed Nov. 18, 2008 in 9 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Order Re Claim Construction, filed Nov. 16, 2009 in 1 page.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Plaintiff Google's Reply to Counterclaims, filed Dec. 8, 2008 in 4 pages.
U.S. District Court Northern District of California, Case No. CV08 04144, *Google Inc. v. Netlist, Inc.*, Stipulation Re: Additional Agreed-Upon Claim Constructions, filed Oct. 28, 2009 in 3 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Complaint for Patent Infringement, filed Dec. 4, 2009 in 47 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Defendant Google Inc.'s Responsive Claim Construction Brief, filed Aug. 4, 2010 in 27 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Exhibit A to Joint Claim Construction and Prehearing Statement under Patent L.R. 4-3, filed Jun. 25, 2010 in 2 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Exhibit B to Joint Claim Construction and Prehearing Statement under Patent L.R. 4-3, filed Jun. 25, 2010 in 23 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Google's Answer to Plaintiff's Complaint for Patent Infringement; and Assertion of Counterclaims, filed Feb. 12, 2010 in 13 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Joint Claim Construction and Prehearing Statement Under Patent Local Rule 4-3, filed Jun. 25, 2010 in 5 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Plaintiff Netlist, Inc.'s Reply Claim Construction Brief, filed Aug. 16, 2010 in 17 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Plaintiff Netlist, Inc.'s Reply to Defendant Google Inc.'s Counterclaim, filed Mar. 8, 2010 in 11 pages.
U.S. District Court Northern District of California, Case No. CV09 05718, *Netlist, Inc. vs. Google, Inc.*, Plaintiff Netlist, Inc.'s Opening Claim Construction Brief, filed Jul. 16, 2010 in 29 pages.
Yao, YL. High Density Memory Selection Circuit; IP.com, IPCOM0000078218D, 1972.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,546 for U.S. Pat. No. 7,289,386 filed May 11, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/001,357 for U.S. Pat. No. 7,636,274; filed Jun. 4, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/001,338; for U.S. Pat. No. 7532537; filed Apr. 19, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/001,339 for U.S. Pat. No. 7619912; filed Jun. 8, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,577 for U.S. Pat. No. 7,289,386 filed Oct. 20, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,578 for U.S. Pat. No. 7,619,912 filed Oct. 20, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/000,579 for U.S. Pat. No. 7,619,912 filed Oct. 20, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/001,381 for U.S. Pat. No. 7,532,537; filed Jun. 9, 2010; Owned by Netlist, Inc. and Its Entire Prosecution History.
Request for *Inter Partes* Reexamination; Reexam U.S. Appl. No. 95/001,758 for U.S. Pat. No. 7,864,627; filed Sep. 15, 2011; Owned by Netlist, Inc. and Its Entire Prosecution History.
U.S. Appl. No. 12/981,380, filed Dec. 29, 2010; and its entire prosecution history.
U.S. Appl. No. 13/154,172, filed Jun. 6, 2011; and its entire prosecution history.
U.S. Appl. No. 12/954,492, filed Nov. 24, 2010; and its entire prosecution history.
"Elpida Memory to Speak at Intel's Memory Implementers Forum Roundtable Event", Intel Developer Forum, [Online]. Retrieved from the Internet: <URSL: <http://www.elpida.com/en/news/2004/0218.html>>, (Jun. 14, 2011), 1 pg.
Micron "DDR SDRAM RDIMM, MT36VDDF12872-1GB, MT36VDDF25672-2GB," 2002 Micron Technology, Inc. 20 pages.
Micron "DDR2 SDRAM Registered DIMM (RDIMM)," 2003 Micron Technology, Inc. 18 pages.
Micron "Synchronous DRAM Module MT18LSDT472," 1998, Micron Technology, Inc., 17 pages.
Micron Technical Note, "Decoupling Capacitor Calculations for a DDR Memory Channel," 2004, 3 pages.
PC133 SDRAM Registered DIMM Design Specification, Revision 1.1, Aug. 1999, 62 pages.
Texas Instruments, "TM2SR72EPN 2097152 by 72-Bit, TM4SRT2EPN 4194304 by 72-Bit, Synchronous Dynamic RAM Modules," 1997, 15 pages.
US District Court Civil Docket; *Google Inc. v. Netlist Inc.*; 4:08cv04144; Date filed Aug. 29, 2008.
US District Court Civil Docket; *Netlist Inc. v. Google Inc.*; 4:09cv5718; Date filed Dec. 4 2009.
US District Court Civil Docket; *Netlist Inc. v. Inphi Corporation*; 2:09cv6900; Date filed Sep. 22, 2002.
Vogt, Pete. "Fully Buffered DIMM (FB-DIMM) Server Memory Architecture: Capacity, Performance, Reliability, and Longevity," Intel, Feb. 18, 2004, 33 pages.
US 6,438,062, 08/2002, Curtis et al. (withdrawn)

* cited by examiner

U.S. Patent

Dec. 20, 2011

Sheet 1 of 23

US 8,081,536 B1

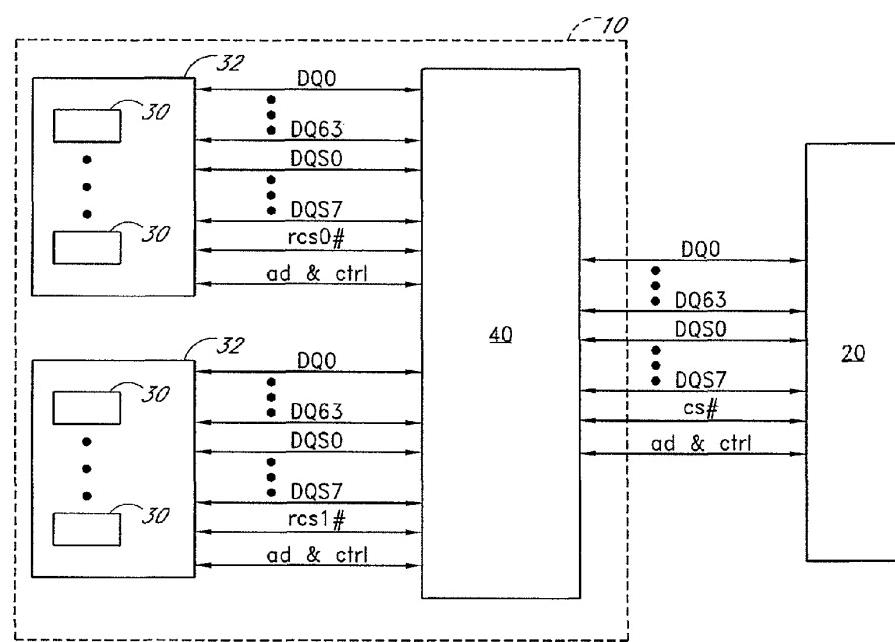


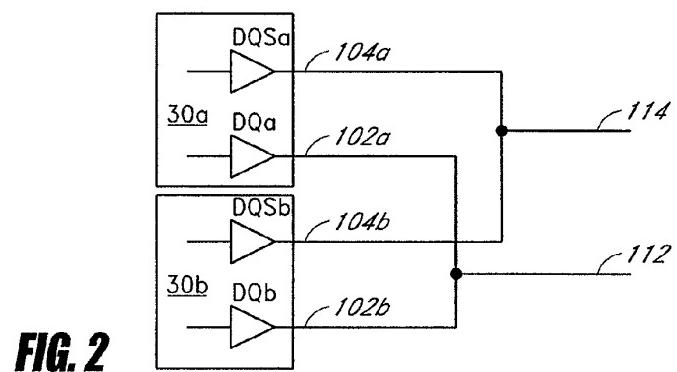
FIG. 1

U.S. Patent

Dec. 20, 2011

Sheet 2 of 23

US 8,081,536 B1

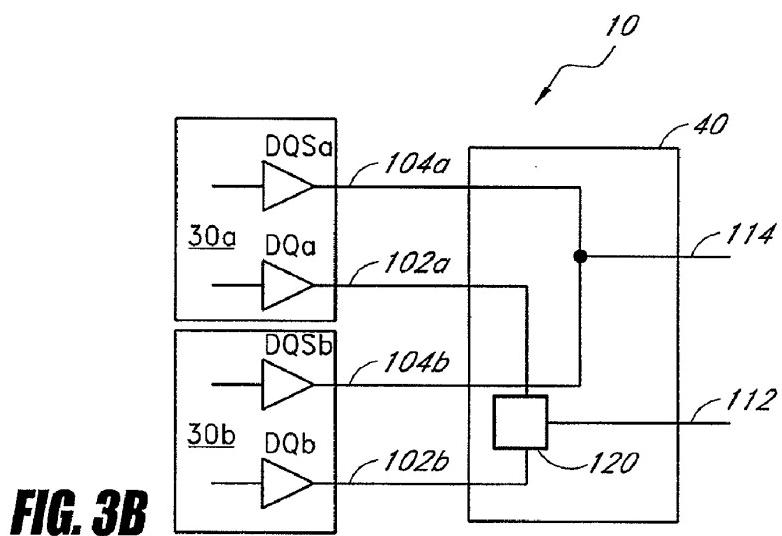
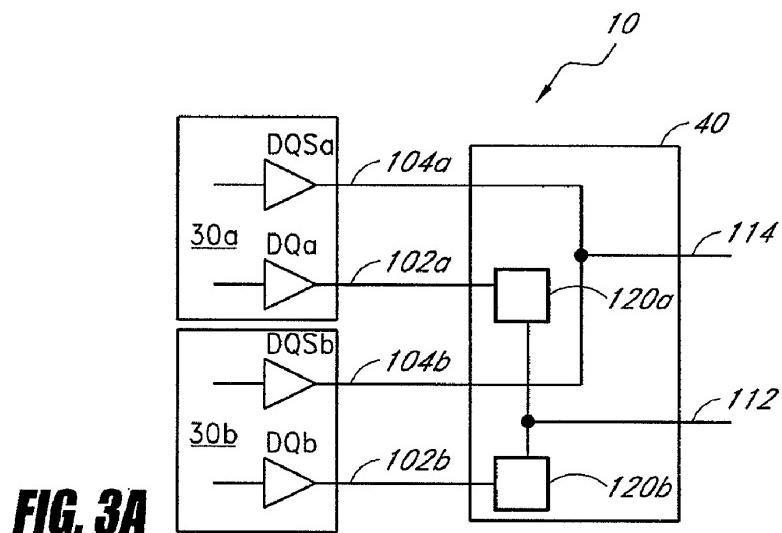


U.S. Patent

Dec. 20, 2011

Sheet 3 of 23

US 8,081,536 B1

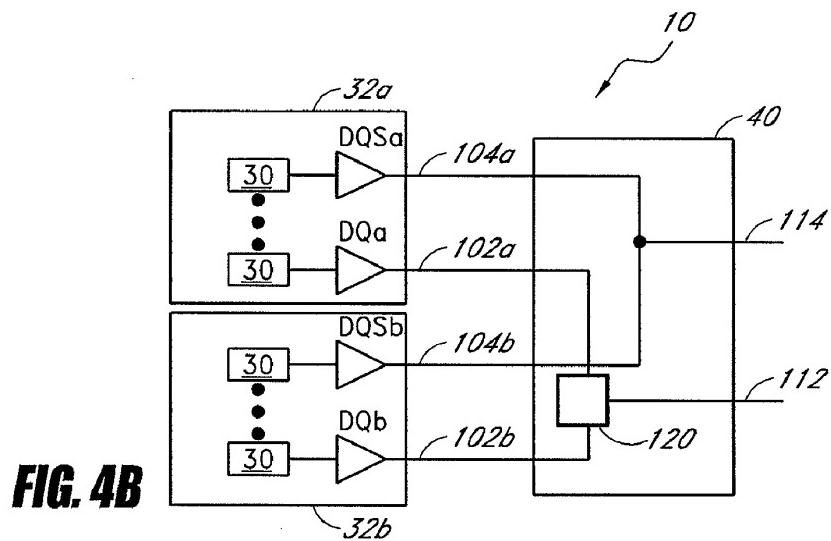
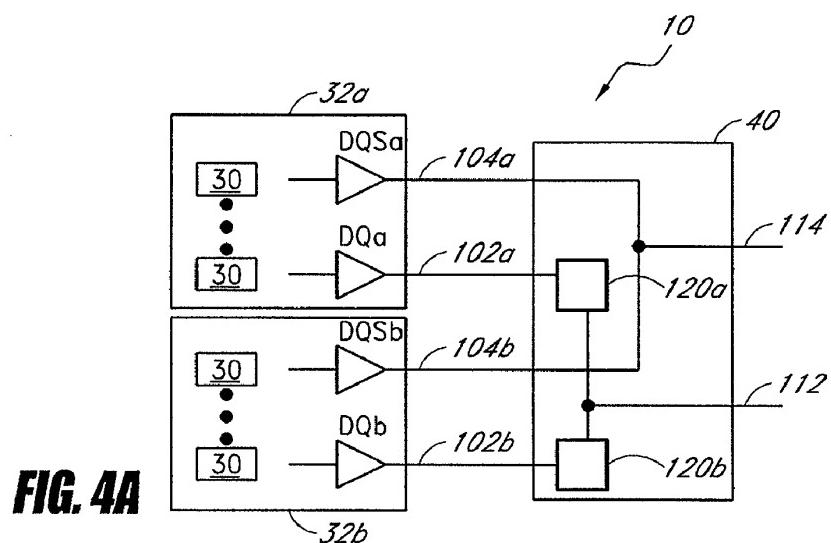


U.S. Patent

Dec. 20, 2011

Sheet 4 of 23

US 8,081,536 B1

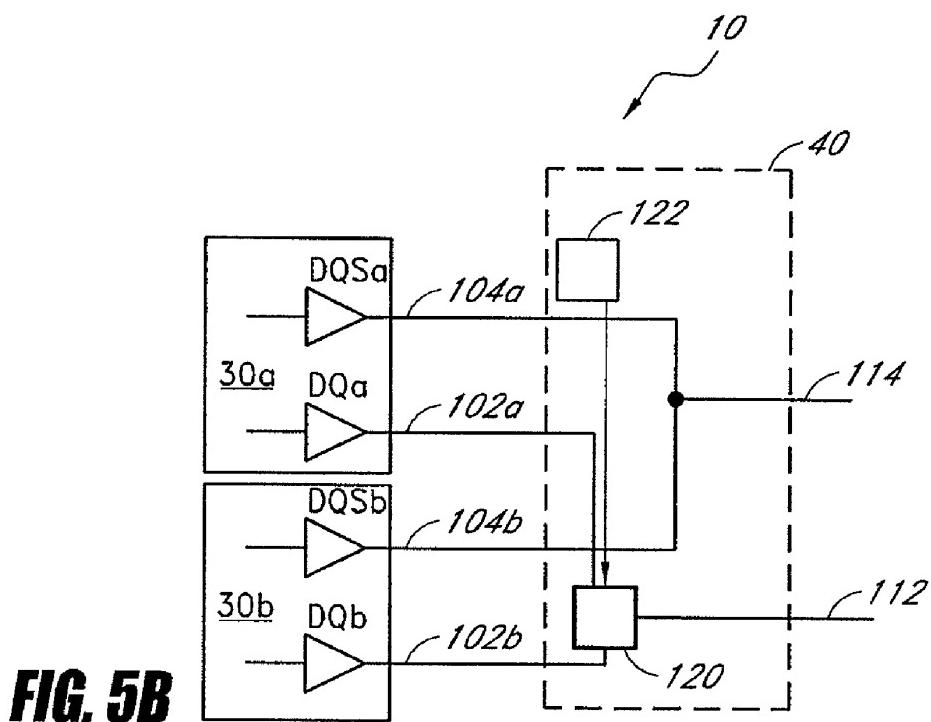
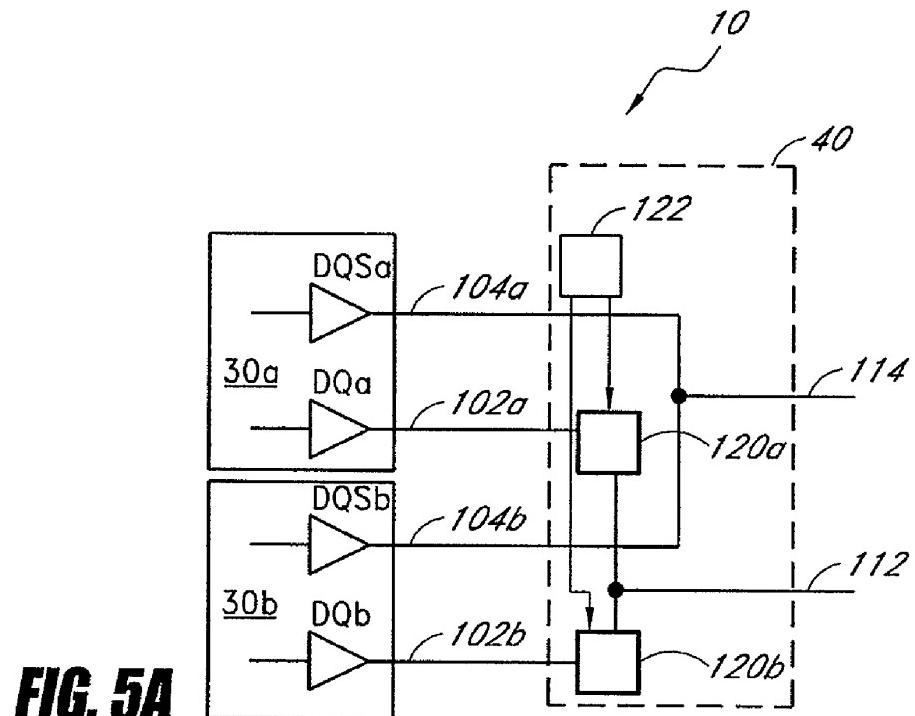


U.S. Patent

Dec. 20, 2011

Sheet 5 of 23

US 8,081,536 B1

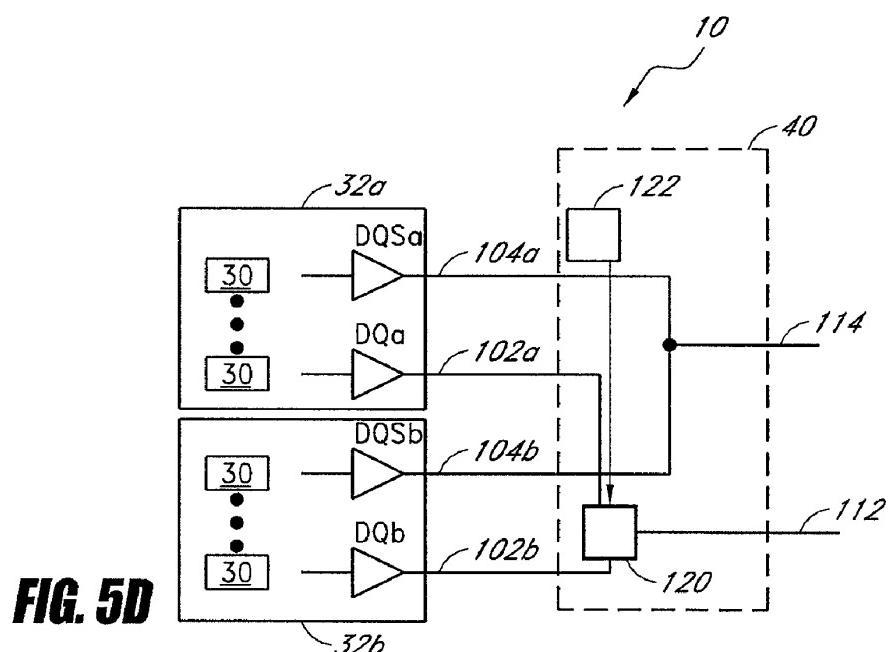
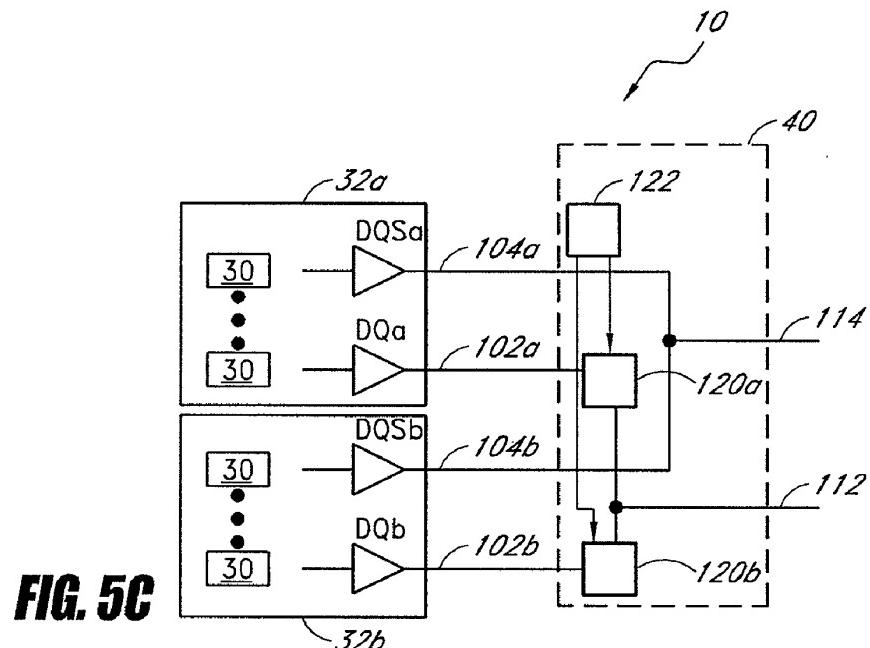


U.S. Patent

Dec. 20, 2011

Sheet 6 of 23

US 8,081,536 B1



U.S. Patent

Dec. 20, 2011

Sheet 7 of 23

US 8,081,536 B1

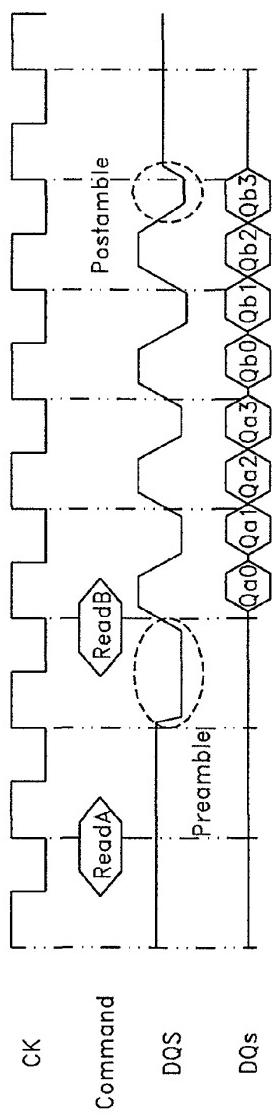


FIG. 6A

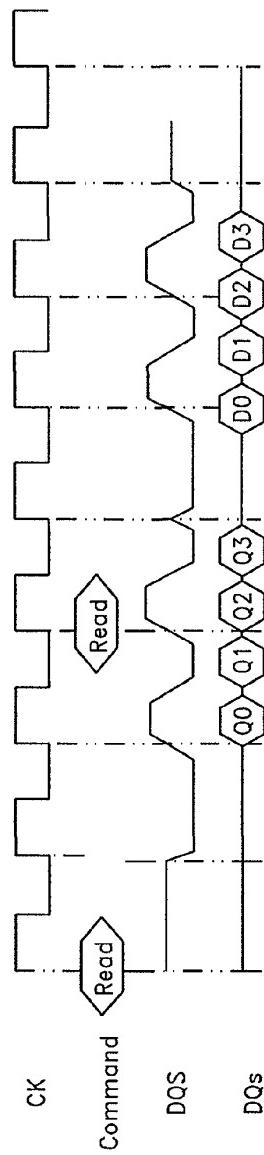


FIG. 6B

U.S. Patent

Dec. 20, 2011

Sheet 8 of 23

US 8,081,536 B1

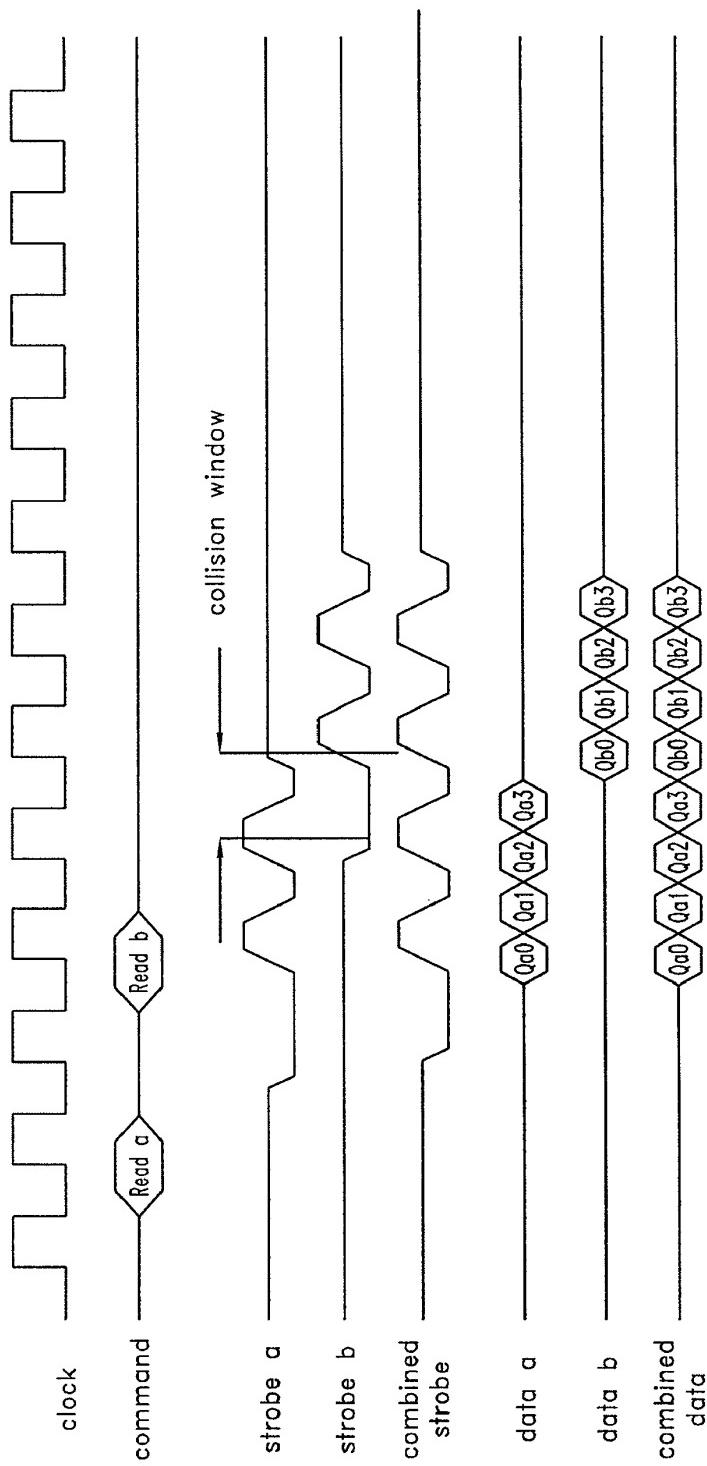


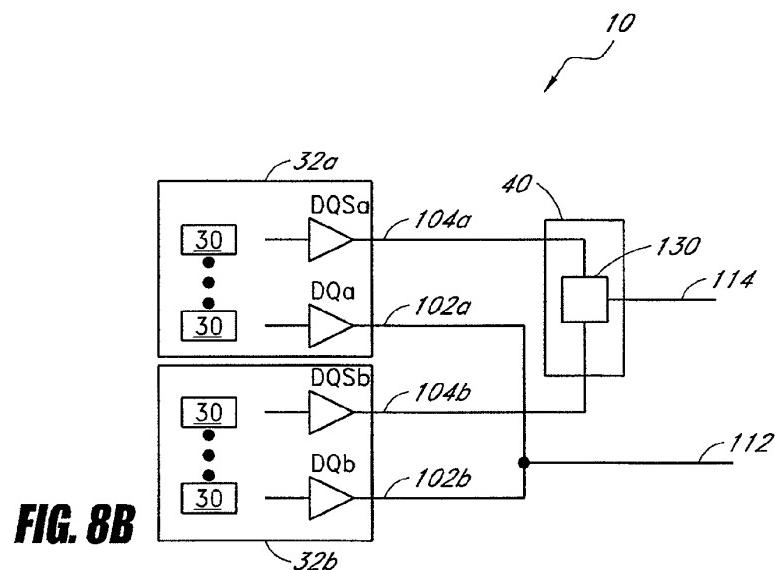
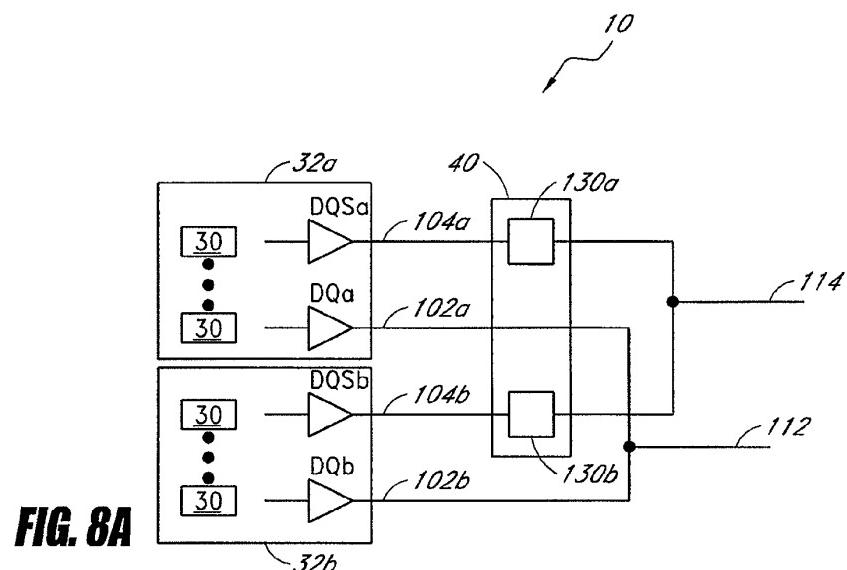
FIG. 7

U.S. Patent

Dec. 20, 2011

Sheet 9 of 23

US 8,081,536 B1

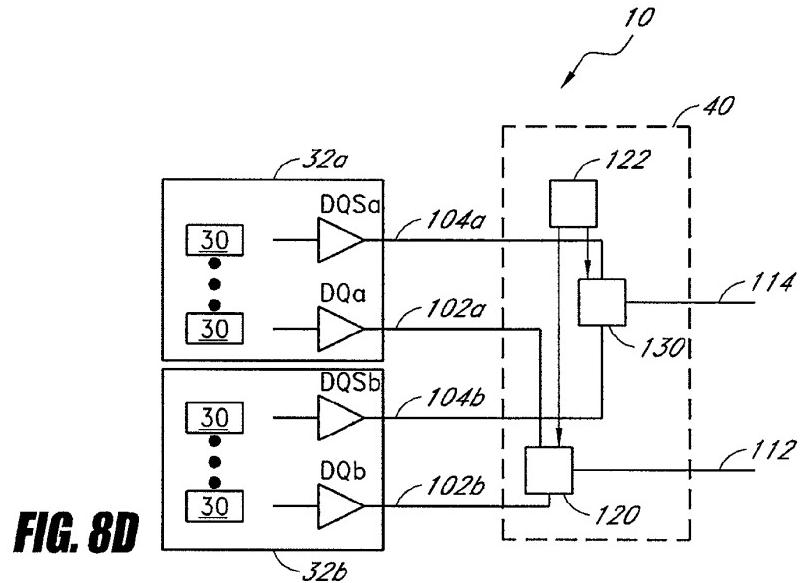
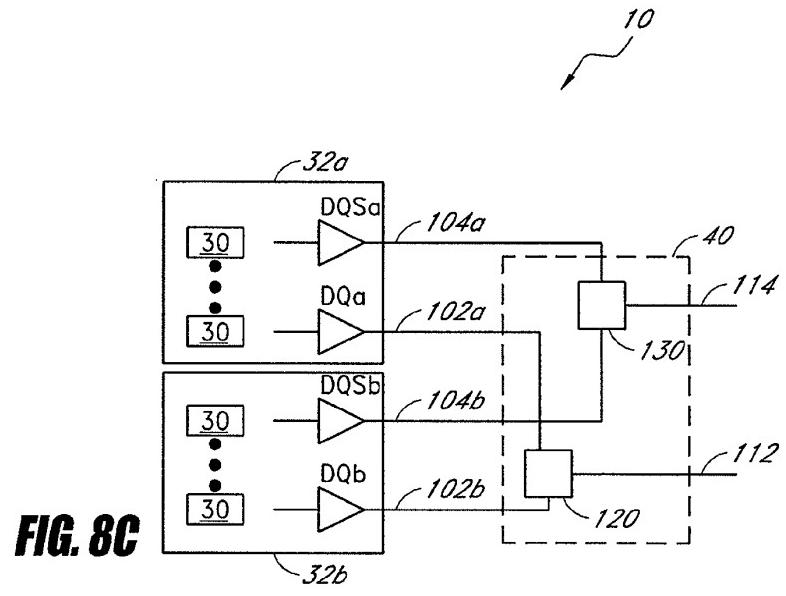


U.S. Patent

Dec. 20, 2011

Sheet 10 of 23

US 8,081,536 B1



U.S. Patent

Dec. 20, 2011

Sheet 11 of 23

US 8,081,536 B1

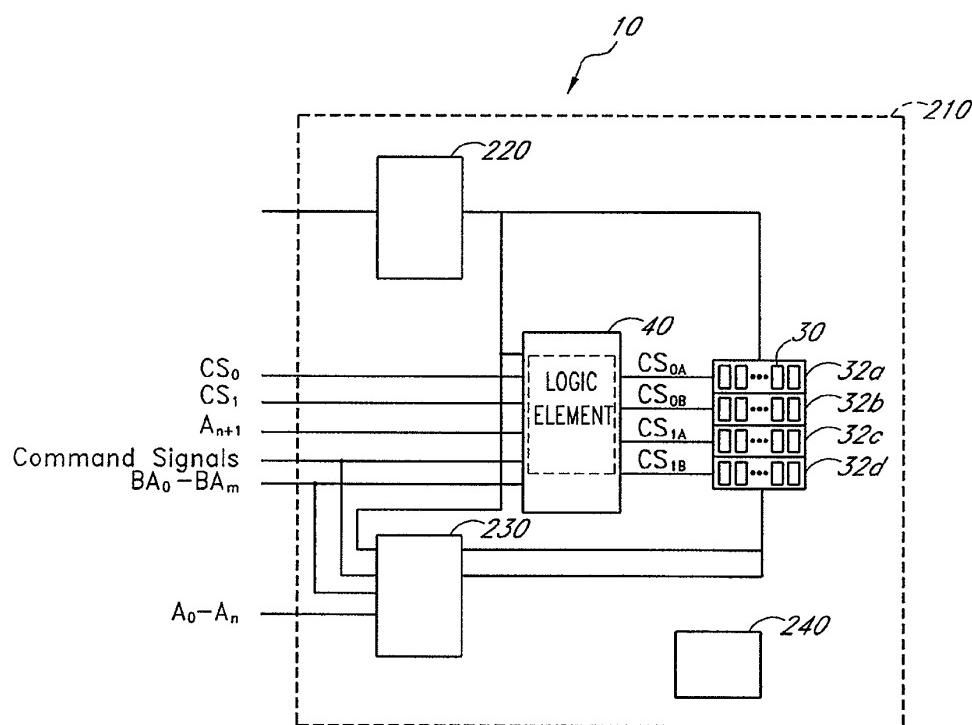


FIG. 9A

U.S. Patent

Dec. 20, 2011

Sheet 12 of 23

US 8,081,536 B1

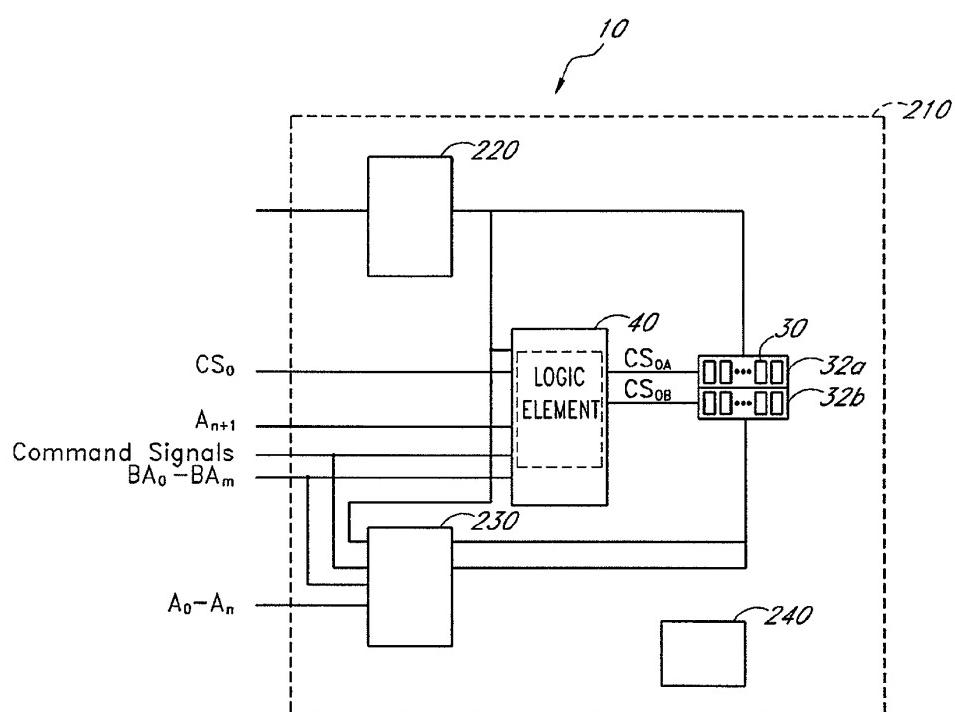


FIG. 9B

U.S. Patent

Dec. 20, 2011

Sheet 13 of 23

US 8,081,536 B1

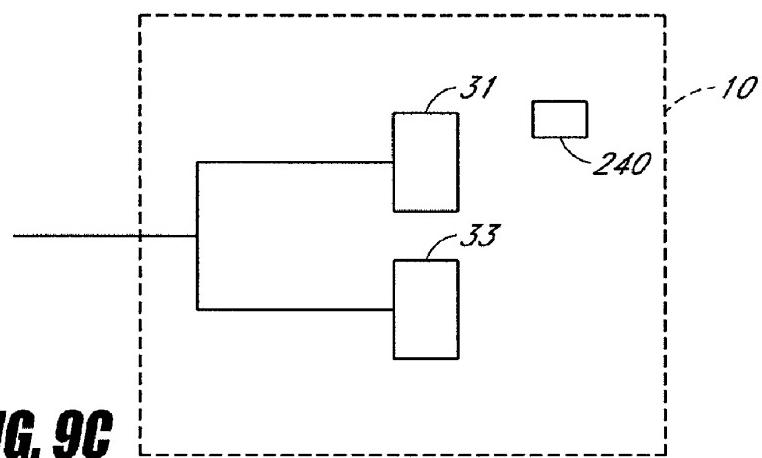


FIG. 9C

U.S. Patent

Dec. 20, 2011

Sheet 14 of 23

US 8,081,536 B1

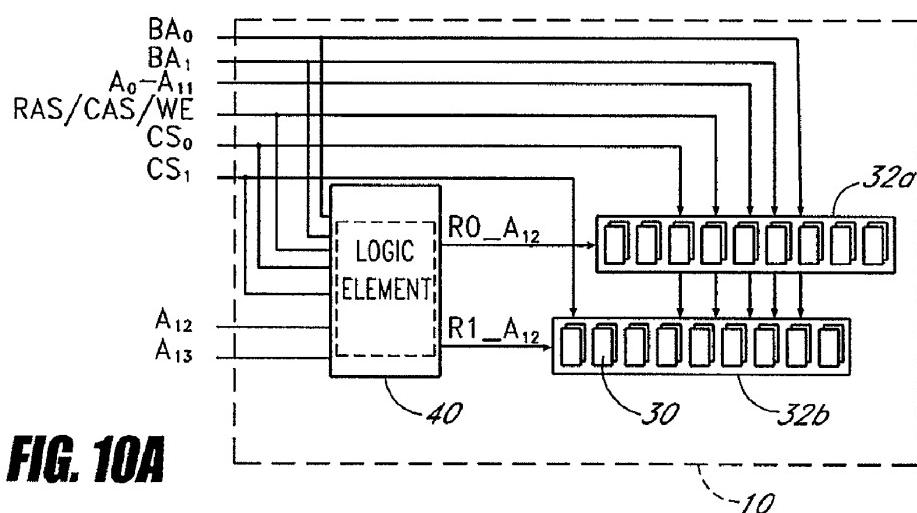


FIG. 10A

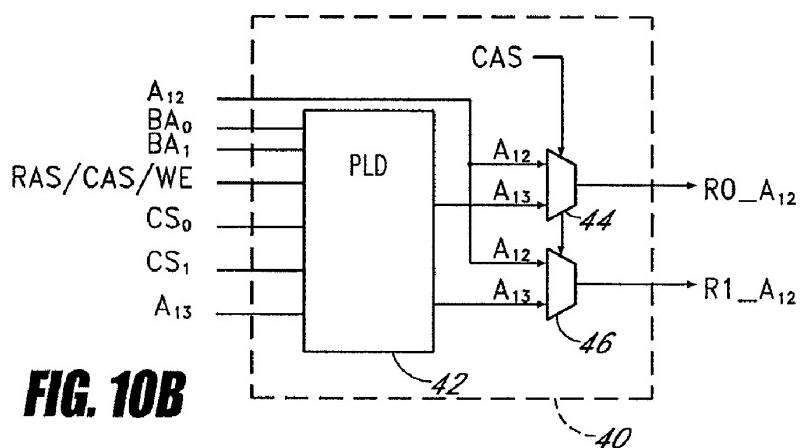


FIG. 10B

U.S. Patent

Dec. 20, 2011

Sheet 15 of 23

US 8,081,536 B1

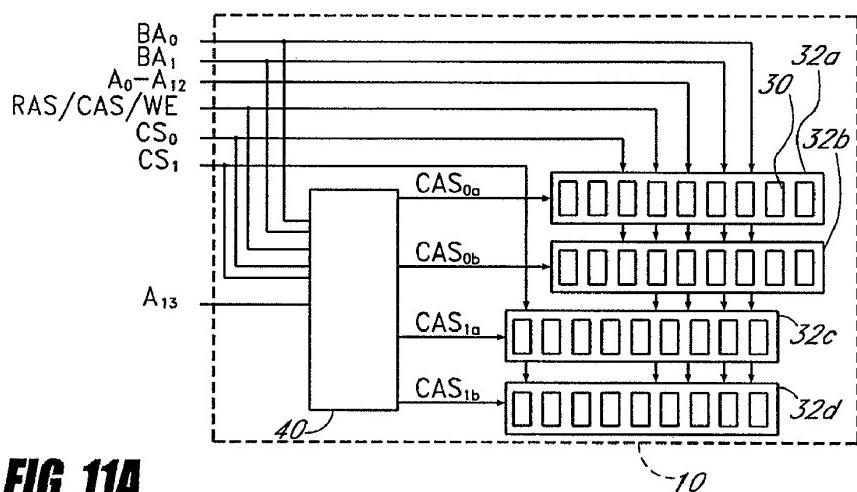


FIG. 11A

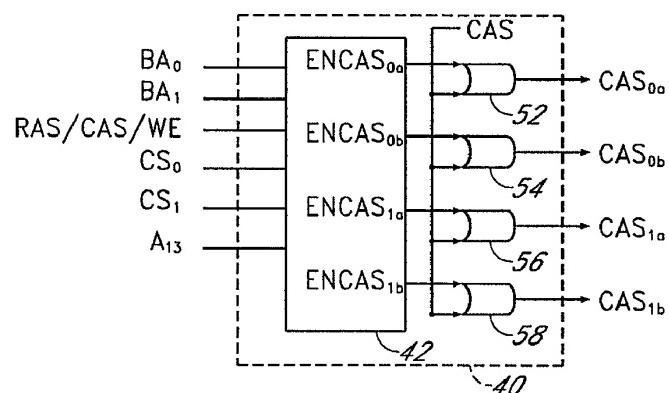


FIG. 11B

U.S. Patent

Dec. 20, 2011

Sheet 16 of 23

US 8,081,536 B1

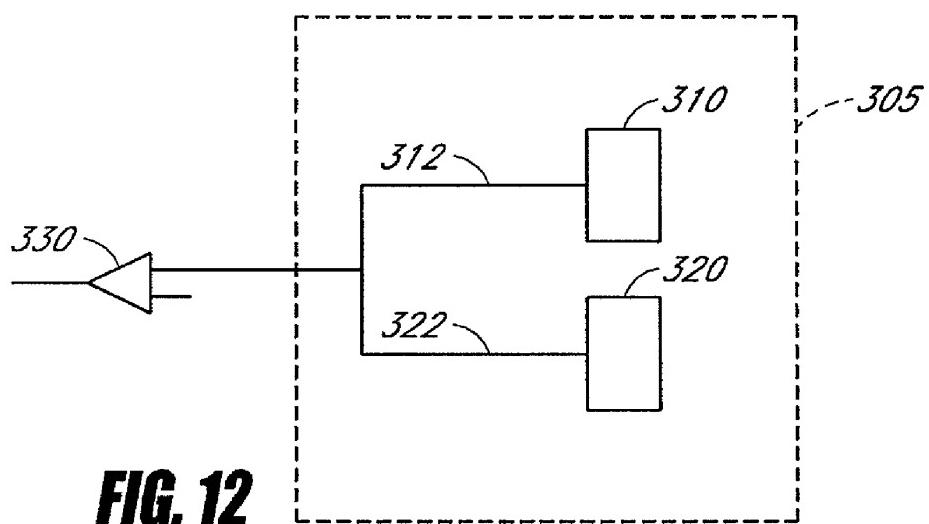


FIG. 12

U.S. Patent

Dec. 20, 2011

Sheet 17 of 23

US 8,081,536 B1

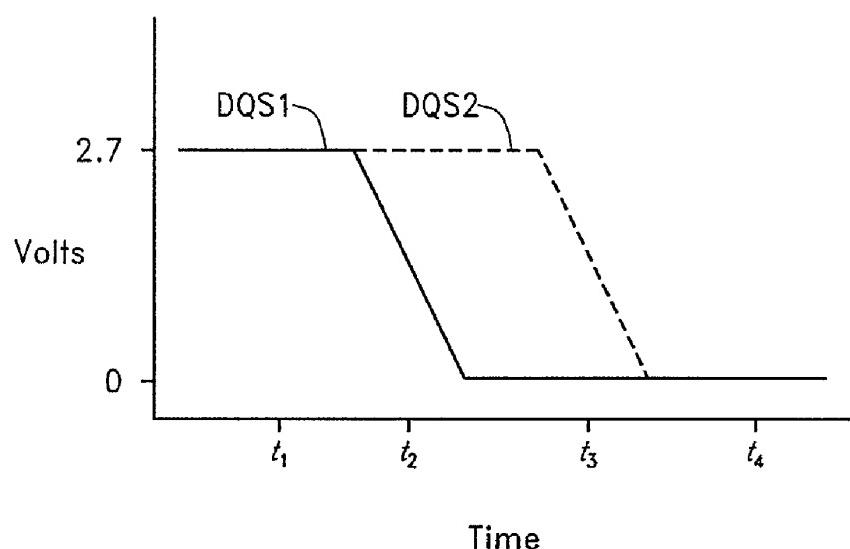


FIG. 13

U.S. Patent

Dec. 20, 2011

Sheet 18 of 23

US 8,081,536 B1

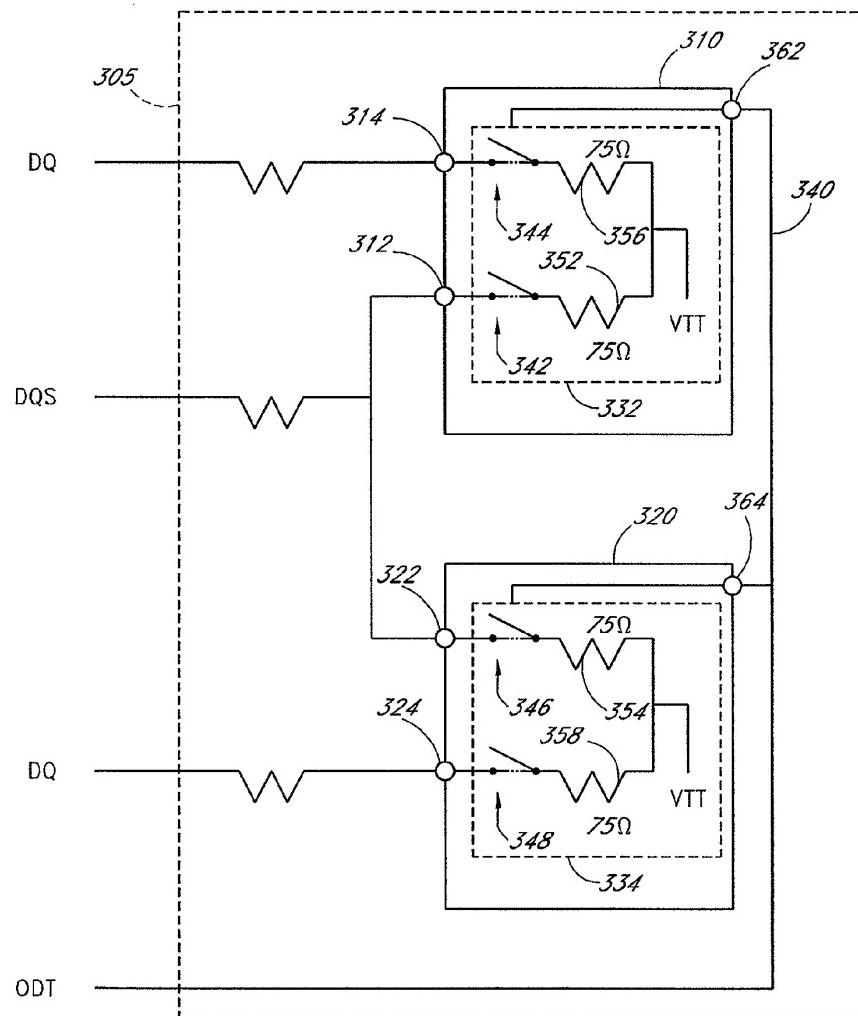


FIG. 14

U.S. Patent

Dec. 20, 2011

Sheet 19 of 23

US 8,081,536 B1

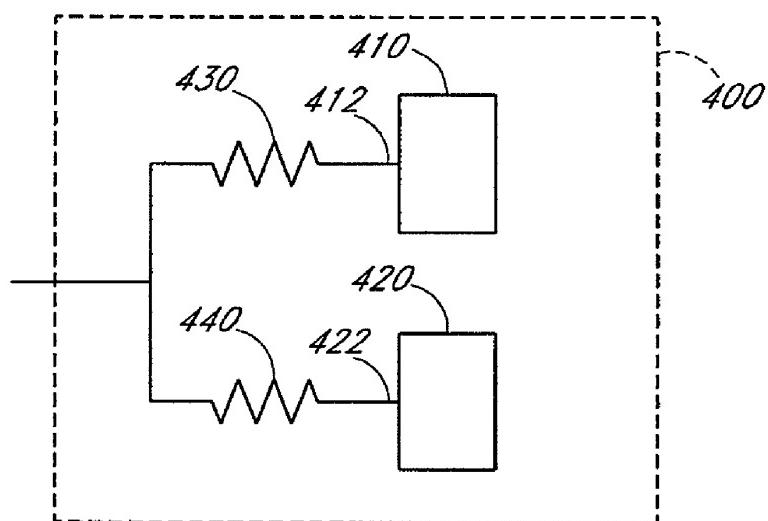


FIG. 15

U.S. Patent

Dec. 20, 2011

Sheet 20 of 23

US 8,081,536 B1

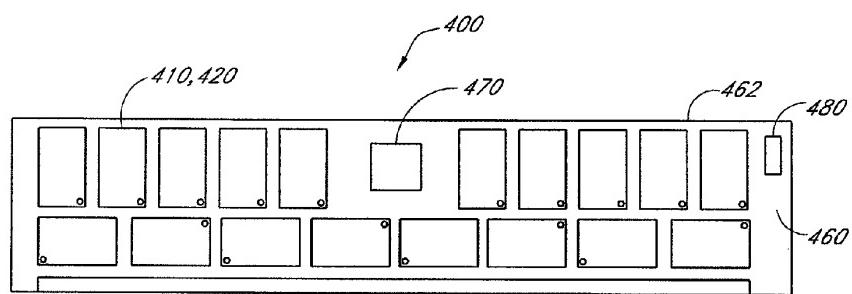


FIG. 16A

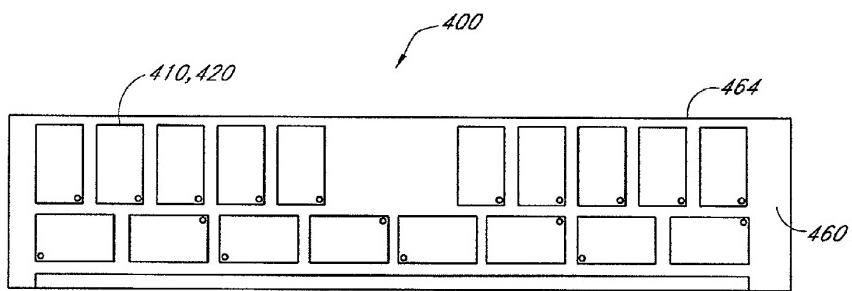


FIG. 16B

U.S. Patent

Dec. 20, 2011

Sheet 21 of 23

US 8,081,536 B1

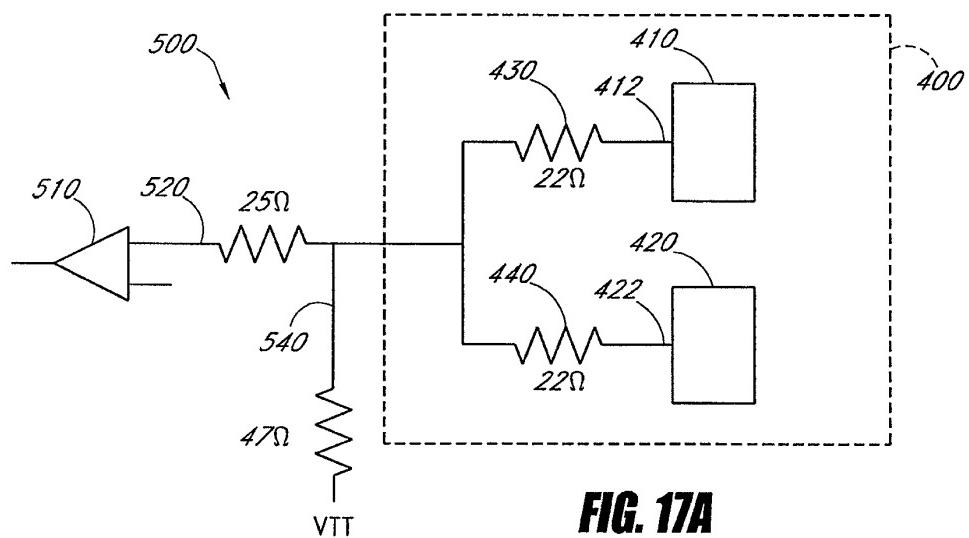


FIG. 17A

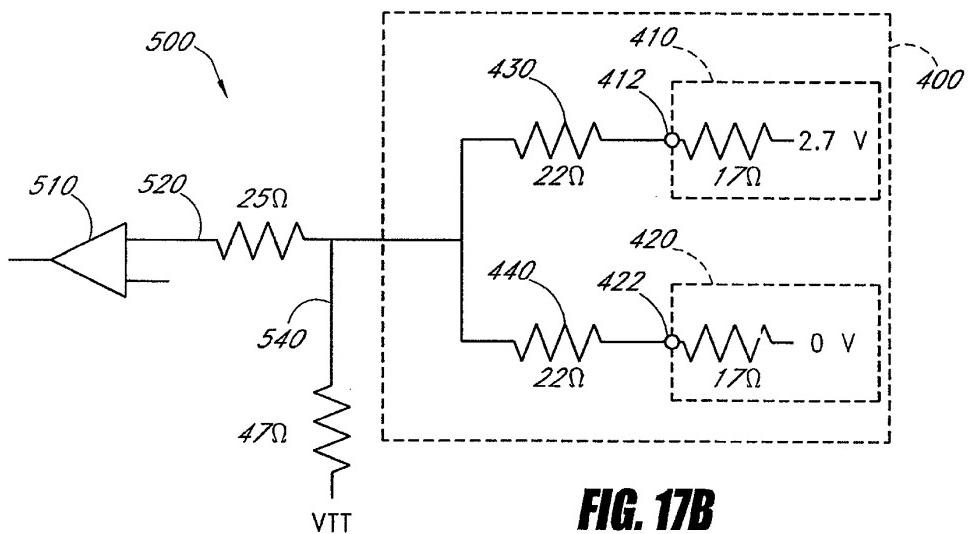


FIG. 17B

U.S. Patent

Dec. 20, 2011

Sheet 22 of 23

US 8,081,536 B1

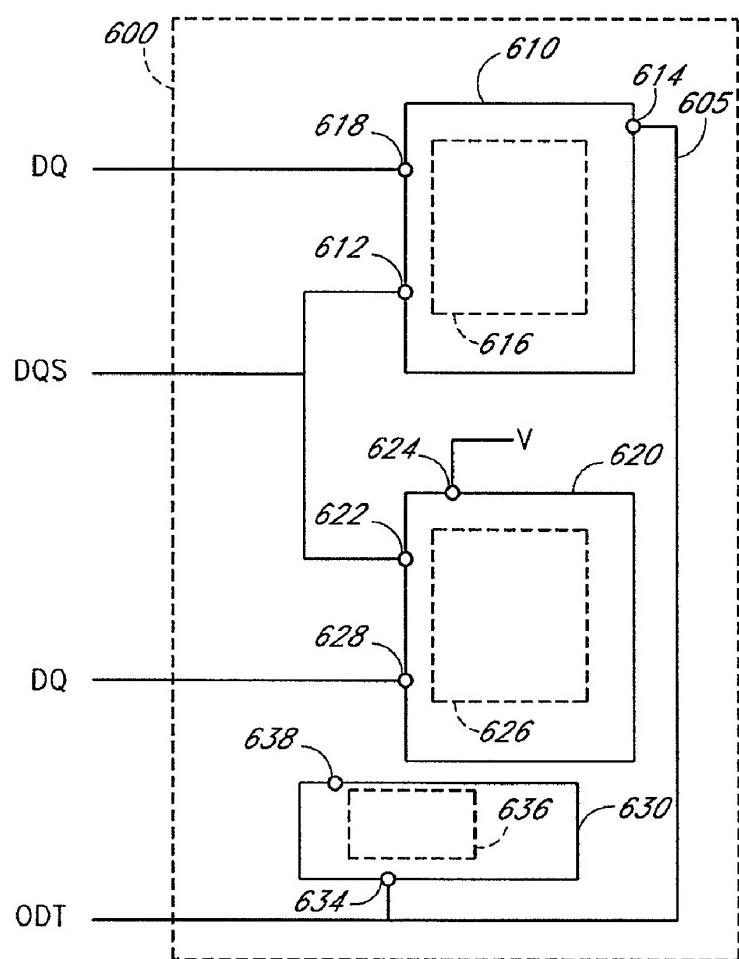


FIG. 18

U.S. Patent

Dec. 20, 2011

Sheet 23 of 23

US 8,081,536 B1

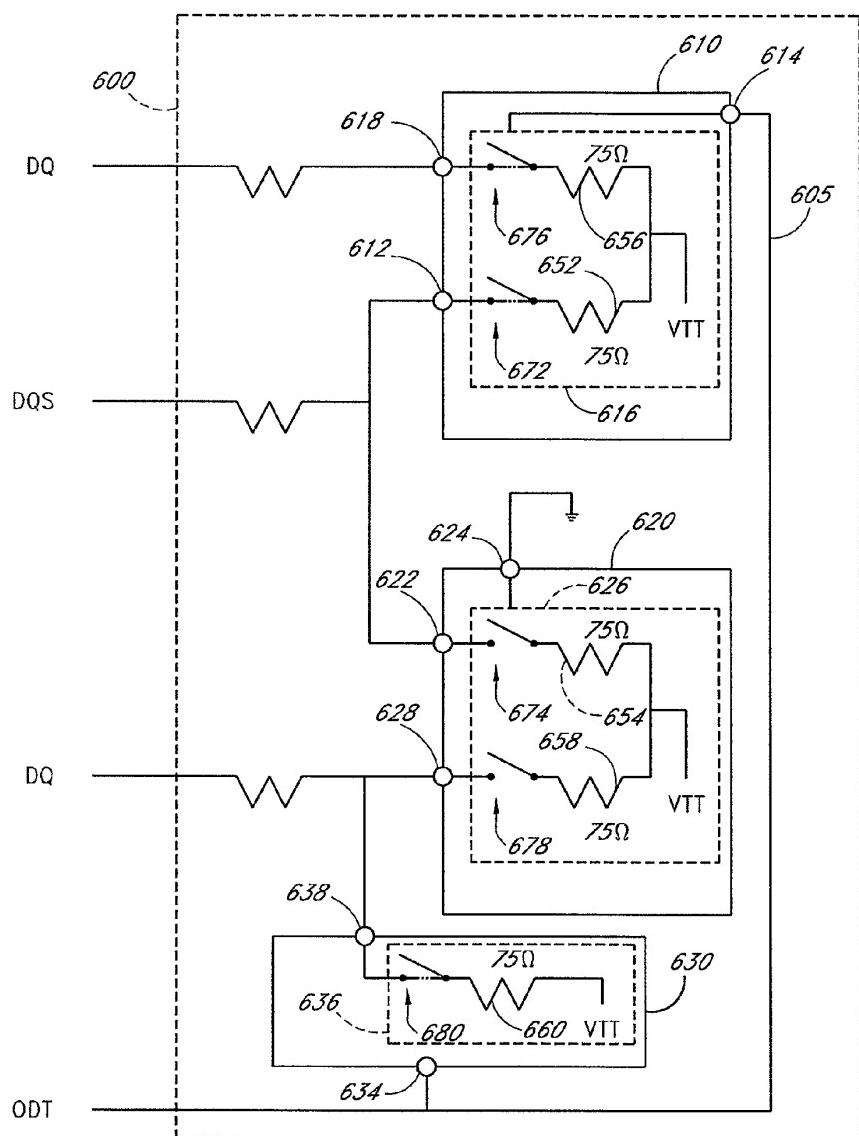


FIG. 19

US 8,081,536 B1

1**CIRCUIT FOR MEMORY MODULE****CROSS-REFERENCE TO RELATED APPLICATIONS**

The present application is a continuation of U.S. patent application Ser. No. 12/995,711, filed Nov. 29, 2010, which is a continuation of U.S. patent application Ser. No. 12/629,827, filed Dec. 2, 2009 and issued as U.S. Pat. No. 7,881,150, which is a continuation of U.S. patent application Ser. No. 12/408,652, filed Mar. 20, 2009 and issued as U.S. Pat. No. 7,636,274, which is a continuation of U.S. patent application Ser. No. 11/335,875, filed Jan. 19, 2006 and issued as U.S. Pat. No. 7,532,537, which claims the benefit of U.S. Provisional Appl. No. 60/645,087, filed Jan. 19, 2005 and which is a continuation-in-part of U.S. patent application Ser. No. 11/173,175, filed Jul. 1, 2005 and issued as U.S. Pat. No. 7,289,386, which claims the benefit of U.S. Provisional Appl. No. 60/588,244, filed Jul. 15, 2004 and which is a continuation-in-part of U.S. patent application Ser. No. 11/075,395, filed Mar. 7, 2005 and issued as U.S. Pat. No. 7,286,436, which claims the benefit of U.S. Provisional Appl. No. 60/550,668, filed Mar. 5, 2004, U.S. Provisional Appl. No. 60/575,595, filed May 28, 2004, and U.S. Provisional Appl. No. 60/590,038, filed Jul. 21, 2004. U.S. patent application Ser. Nos. 12/995,711, 12/629,827, 12/408,652, 11/335,875, 11/173,175, and 11/075,395 and U.S. Provisional Appl. Nos. 60/550,668, 60/575,595, 60/590,038, 60/588,244, and 60/645,087 are each incorporated in their entirety by reference herein.

BACKGROUND OF THE INVENTION**1. Field of the Invention**

The present invention relates generally to memory modules of a computer system, and more specifically to devices and methods for improving the performance, the memory capacity, or both, of memory modules.

2. Description of the Related Art

Certain types of memory modules comprise a plurality of dynamic random-access memory (DRAM) devices mounted on a printed circuit board (PCB). These memory modules are typically mounted in a memory slot or socket of a computer system (e.g., a server system or a personal computer) and are accessed by the processor of the computer system. Memory modules typically have a memory configuration with a unique combination of rows, columns, and banks which result in a total memory capacity for the memory module.

For example, a 512-Megabyte memory module (termed a "512-MB" memory module, which actually has 2^{29} or 536,870,912 bytes of capacity) will typically utilize eight 512-Megabit DRAM devices (each identified as a "512-Mb" DRAM device, each actually having 2^{29} or 536,870,912 bits of capacity). The memory cells (or memory locations) of each 512-Mb DRAM device can be arranged in four banks, with each bank having an array of 2^{24} (or 16,777,216) memory locations arranged as 2^{13} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 64M 8-bit-wide memory locations (actually with four banks of 2^{27} or 134,217,728 one-bit memory cells arranged to provide a total of 2^{26} or 67,108,864 memory locations with 8 bits each) are identified as having a "64 Mb×8" or "64M×8-bit" configuration, or as having a depth of 64M and a bit width of 8. Furthermore, certain commercially-available 512-MB memory modules are termed to have a "64M×8-byte" configuration or a "64M×64-bit" configuration with a depth of 64M and a width of 8 bytes or 64 bits.

Similarly, a 1-Gigabyte memory module (termed a "1-GB" memory module, which actually has 2^{30} or 1,073,741,824 bytes of capacity) can utilize eight 1-Gigabit DRAM devices (each identified as a "1-Gb" DRAM device, each actually having 2^{30} or 1,073,741,824 bits of capacity). The memory locations of each 1-Gb DRAM device can be arranged in four banks, with each bank having an array of memory locations with 2^{14} rows and 2^{11} columns, and with each memory location having a width of 8 bits. Such DRAM devices with 128M 8-bit-wide memory locations (actually with a total of 2^{27} or 134,217,728 memory locations with 8 bits each) are identified as having a "128 Mb×8" or "128M×8-bit" configuration, or as having a depth of 128M and a bit width of 8. Furthermore, certain commercially-available 1-GB memory modules are identified as having a "128M×8-byte" configuration or a "128M×64-bit" configuration with a depth of 128M and a width of 8 bytes or 64 bits.

The commercially-available 512-MB (64M×8-byte) memory modules and the 1-GB (128M×8-byte) memory modules described above are typically used in computer systems (e.g., personal computers) which perform graphics applications since such "×8" configurations are compatible with data mask capabilities often used in such graphics applications. Conversely, memory modules with "×4" configurations are typically used in computer systems such as servers which are not as graphics-intensive. Examples of such commercially available "×4" memory modules include, but are not limited to, 512-MB (128M×4-byte) memory modules comprising eight 512-Mb (128 Mb×4) memory devices.

The DRAM devices of a memory module are generally arranged as ranks or rows of memory, each rank of memory generally having a bit width. For example, a memory module in which each rank of the memory module is 64 bits wide is described as having an "x64" organization. Similarly, a memory module having 72-bit-wide ranks is described as having an "x72" organization.

The memory capacity of a memory module increases with the number of memory devices. The number of memory devices of a memory module can be increased by increasing the number of memory devices per rank or by increasing the number of ranks. For example, a memory module with four ranks has double the memory capacity of a memory module with two ranks and four times the memory capacity of a memory module with one rank. Rather than referring to the memory capacity of the memory module, in certain circumstances, the memory density of the memory module is referred to instead.

During operation, the ranks of a memory module are selected or activated by address and command signals that are received from the processor. Examples of such address and command signals include, but are not limited to, rank-select signals, also called chip-select signals. Most computer and server systems support one-rank and two-rank memory modules. By only supporting one-rank and two-rank memory modules, the memory density that can be incorporated in each memory slot is limited.

Various aspects of the design of a memory module impose limitations on the size of the memory arrays of the memory module. Certain such aspects are particularly important for memory modules designed to operate at higher frequencies. Examples of such aspects include, but are not limited to, memory device (e.g., chip) densities, load fan-out, signal integrity, available rank selects, power dissipation, and thermal profiles.

SUMMARY OF THE INVENTION

In certain embodiments, a circuit is configured to be mounted on a memory module configured to be operationally

US 8,081,536 B1

3

coupled to a computer system. The memory modules has a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of chip-select signals. The circuit is configurable to receive a set of signals comprising address signals and a second number of chip-select signals smaller than the first number of chip-select signals. The circuit is further configurable to generate phase-locked clock signals, wherein the circuit is configurable to transmit the phase-locked clock signals to the first number of ranks. The circuit is further configurable to selectively isolate a load of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals. The circuit is further configurable to generate the first number of chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of chip-select signals.

In certain embodiments, a method operates a memory module configured to be operationally coupled to a computer system. The memory module has a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of chip-select signals. The method comprises receiving a set of signals comprising address signals and a second number of chip-select signals smaller than the first number of chip-select signals. The method further comprises using the memory module to generate phase-locked clock signals and transmitting the phase-locked clock signals to the first number of ranks. The method further comprises selectively isolating a load of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals. The method further comprises generating the first number of chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of chip-select signals.

In certain embodiments, a circuit is configured to be mounted on a memory module configured to be operationally coupled to a computer system. The memory module has a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of chip-select signals. The circuit is configurable to receive a set of signals comprising a command signal, address signals, and a second number of chip-select signals smaller than the first number of chip-select signals. The circuit is further configurable to generate phase-locked clock signals, wherein the circuit is configurable to transmit the phase-locked clock signals to the first number of ranks. The circuit is further configurable to store an address signal of the set of signals during a row access procedure and to pass the stored address signal as an address signal to the first number of ranks during a subsequent column address procedure. The circuit is further configurable to respond at least in part to the phase-locked clock signals, address signals, and the second number of chip-select signals by selecting at least one rank of the first number of ranks to receive the command signal and transmitting the command signal to the selected at least one rank. The circuit is further configurable to selectively isolate a data signal line load of at least one rank of the first number of ranks from the computer system in response at least in part to the set

4

of signals by presenting a load of the circuit to the computer system. The circuit is further configurable to generate the first number of chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of chip-select signals.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 schematically illustrates an example memory module in accordance with certain embodiments described herein.

FIG. 2 schematically illustrates a circuit diagram of two memory devices of a conventional memory module.

FIGS. 3A and 3B schematically illustrate example memory modules having a circuit which selectively isolates one or both of the DQ data signal lines of the two memory devices from the computer system in accordance with certain embodiments described herein.

FIGS. 4A and 4B schematically illustrate example memory modules having a circuit which selectively isolates one or both of the DQ data signal lines of the two ranks of memory devices from the computer system in accordance with certain embodiments described herein.

FIGS. 5A-5D schematically illustrate example memory modules having a circuit comprising a logic element and one or more switches operatively coupled to the logic element in accordance with certain embodiments described herein.

FIG. 6A shows an exemplary timing diagram of a gapless read burst for a back-to-back adjacent read condition from one memory device.

FIG. 6B shows an exemplary timing diagram with an extra clock cycle between successive read commands issued to different memory devices for successive read accesses from different memory devices.

FIG. 7 shows an exemplary timing diagram in which the last data strobe of memory device "a" collides with the preamble time interval of the data strobe of memory device "b."

FIGS. 8A-8D schematically illustrate circuit diagrams of example memory modules comprising a circuit which multiplexes the DQS data strobe signal lines from one another in accordance with certain embodiments described herein.

FIG. 9A schematically illustrates an example memory module with four ranks of memory devices compatible with certain embodiments described herein.

FIG. 9B schematically illustrates an example memory module with two ranks of memory devices compatible with certain embodiments described herein.

FIG. 9C schematically illustrates another example memory module in accordance with certain embodiments described herein.

FIG. 10A schematically illustrates an exemplary memory module which doubles the rank density in accordance with certain embodiments described herein.

FIG. 10B schematically illustrates an exemplary circuit compatible with embodiments described herein.

FIG. 11A schematically illustrates an exemplary memory module which doubles number of ranks in accordance with certain embodiments described herein.

FIG. 11B schematically illustrates an exemplary circuit compatible with embodiments described herein.

FIG. 12 schematically illustrates an exemplary memory module in which a data strobe (DQS) pin of a first memory device is electrically connected to a DQS pin of a second memory device while both DQS pins are active.

FIG. 13 is an exemplary timing diagram of the voltages applied to the two DQS pins due to non-simultaneous switching.

US 8,081,536 B1

5

FIG. 14 schematically illustrates another exemplary memory module in which a DQS pin of a first memory device is connected to a DQS pin of a second memory device.

FIG. 15 schematically illustrates an exemplary memory module in accordance with certain embodiments described herein.

FIGS. 16A and 16B schematically illustrate a first side and a second side, respectively, of a memory module with eighteen 64Mx4 bit, DDR-1 SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board.

FIGS. 17A and 17B schematically illustrate an exemplary embodiment of a memory module in which a first resistor and a second resistor are used to reduce the current flow between the first DQS pin and the second DQS pin.

FIG. 18 schematically illustrates another exemplary memory module compatible with certain embodiments described herein.

FIG. 19 schematically illustrates a particular embodiment of the memory module schematically illustrated by FIG. 18.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

Load Isolation

FIG. 1 schematically illustrates an example memory module 10 compatible with certain embodiments described herein. The memory module 10 is connectable to a memory controller 20 of a computer system (not shown). The memory module 10 comprises a plurality of memory devices 30, each memory device 30 having a corresponding load. The memory module 10 further comprises a circuit 40 electrically coupled to the plurality of memory devices 30 and configured to be electrically coupled to the memory controller 20 of the computer system. The circuit 40 selectively isolates one or more of the loads of the memory devices from the computer system. The circuit 40 comprises logic which translates between a system memory domain of the computer system and a physical memory domain of the memory module 10.

As used herein, the term "load" is a broad term which includes, without limitation, electrical load, such as capacitive load, inductive load, or impedance load. As used herein, the term "isolation" is a broad term which includes, without limitation, electrical separation of one or more components from another component or from one another. As used herein, the term "circuit" is a broad term which includes, without limitation, an electrical component or device, or a configuration of electrical components or devices which are electrically or electromagnetically coupled together (e.g., integrated circuits), to perform specific functions.

Various types of memory modules 10 are compatible with embodiments described herein. For example, memory modules 10 having memory capacities of 512-MB, 1-GB, 2-GB, 4-GB, 8-GB, as well as other capacities, are compatible with embodiments described herein. Certain embodiments described herein are applicable to various frequencies including, but not limited to 100 MHz, 200 MHz, 400 MHz, 800 MHz, and above. In addition, memory modules 10 having widths of 4 bytes, 8 bytes, 16 bytes, 32 bytes, or 32 bits, 64 bits, 128 bits, 256 bits, as well as other widths (in bytes or in bits), are compatible with embodiments described herein. In certain embodiments, the memory module 10 comprises a printed circuit board on which the memory devices 30 are mounted, a plurality of edge connectors configured to be electrically coupled to a corresponding plurality of contacts of a module slot of the computer system, and a plurality of electrical conduits which electrically couple the memory devices 30 to the circuit 40 and which electrically couple the

6

circuit 40 to the edge connectors. Furthermore, memory modules 10 compatible with embodiments described herein include, but are not limited to, single in-line memory modules (SIMMs), dual in-line memory modules (DIMMs), small-outline DIMMs (SO-DIMMs), unbuffered DIMMs (UDIMMs), registered DIMMs (RDIMMs), fully-buffered DIMM (FBDIMM), rank-buffered DIMMs (RBDIMMs), mini-DIMMs, and micro-DIMMs.

Memory devices 30 compatible with embodiments described herein include, but are not limited to, random-access memory (RAM), dynamic random-access memory (DRAM), synchronous DRAM (SDRAM), and double-data-rate DRAM (e.g., SDR, DDR-1, DDR-2, DDR-3). In addition, memory devices 30 having bit widths of 4, 8, 16, 32, as well as other bit widths, are compatible with embodiments described herein. Memory devices 30 compatible with embodiments described herein have packaging which include, but are not limited to, thin small-outline package (TSOP), ball-grid-array (BGA), fine-pitch BGA (FBGA), micro-BGA (μBGA), mini-BGA (mBGA), and chip-scale packaging (CSP). Memory devices 30 compatible with embodiments described herein are available from a number of sources, including but not limited to, Samsung Semiconductor, Inc. of San Jose, Calif., Infineon Technologies AG of San Jose, Calif., and Micron Technology, Inc. of Boise, Id. Persons skilled in the art can select appropriate memory devices 30 in accordance with certain embodiments described herein.

In certain embodiments, the plurality of memory devices 30 comprises a first number of memory devices 30. In certain such embodiments, the circuit 40 selectively isolates a second number of the memory devices 30 from the computer system, with the second number less than the first number.

In certain embodiments, the plurality of memory devices 30 are arranged in a first number of ranks. For example, in certain embodiments, the memory devices 30 are arranged in two ranks, as schematically illustrated by FIG. 1. In other embodiments, the memory devices 30 are arranged in four ranks. Other numbers of ranks of the memory devices 30 are also compatible with embodiments described herein.

In certain embodiments, the circuit comprises a logic element selected from a group consisting of: a programmable-logic device (PLD), an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a custom-designed semiconductor device, and a complex programmable-logic device (CPLD). In certain embodiments, the logic element of the circuit 40 is a custom device. Sources of logic elements compatible with embodiments described herein include, but are not limited to, Lattice Semiconductor Corporation of Hillsboro, Oreg., Altera Corporation of San Jose, Calif., and Xilinx Incorporated of San Jose, Calif. In certain embodiments, the logic element comprises various discrete electrical elements, while in certain other embodiments, the logic element comprises one or more integrated circuits.

In certain embodiments, the circuit 40 further comprises one or more switches which are operatively coupled to the logic element to receive control signals from the logic element. Examples of switches compatible with certain embodiments described herein include, but are not limited to, field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex.

FIG. 2 schematically illustrates a circuit diagram of two memory devices 30a, 30b of a conventional memory module showing the interconnections between the DQ data signal lines 102a, 102b of the memory devices 30a, 30b and the DQS data strobe signal lines 104a, 104b of the memory

US 8,081,536 B1

7

devices 30a, 30b. Each of the memory devices 30a, 30b has a plurality of DQ data signal lines and a plurality of DQS data strobe signal lines, however, for simplicity, FIG. 2 only illustrates a single DQ data signal line and a single DQS data strobe signal line for each memory device 30a, 30b. The DQ data signal lines 102a, 102b and the DQS data strobe signal lines 104a, 104b are typically conductive traces etched on the printed circuit board of the memory module. As shown in FIG. 2, each of the memory devices 30a, 30b has their DQ data signal lines 102a, 102b electrically coupled to a common DQ line 112 and their DQS data strobe signal lines 104a, 104b electrically coupled to a common DQS line 114. The common DQ line 112 and the common DQS line 114 are electrically coupled to the memory controller 20 of the computer system. Thus, the computer system is exposed to the loads of both memory devices 30a, 30b concurrently.

In certain embodiments, the circuit 40 selectively isolates the loads of at least some of the memory devices 30 from the computer system. The circuit 40 of certain embodiments is configured to present a significantly reduced load to the computer system. In certain embodiments in which the memory devices 30 are arranged in a plurality of ranks, the circuit 40 selectively isolates the loads of some (e.g., one or more) of the ranks of the memory module 10 from the computer system. In certain other embodiments, the circuit 40 selectively isolates the loads of all of the ranks of the memory module 10 from the computer system. For example, when a memory module 10 is not being accessed by the computer system, the capacitive load on the memory controller 20 of the computer system by the memory module 10 can be substantially reduced to the capacitive load of the circuit 40 of the memory module 10.

As schematically illustrated by FIGS. 3A and 3B, an example memory module 10 compatible with certain embodiments described herein comprises a circuit 40 which selectively isolates one or both of the DQ data signal lines 102a, 102b of the two memory devices 30a, 30b from the common DQ data signal line 112 coupled to the computer system. Thus, the circuit 40 selectively allows a DQ data signal to be transmitted from the memory controller 20 of the computer system to one or both of the DQ data signal lines 102a, 102b. In addition, the circuit 40 selectively allows one of a first DQ data signal from the DQ data signal line 102a of the first memory device 30a or a second DQ data signal from the DQ data signal line 102b of the second memory device 30b to be transmitted to the memory controller 20 via the common DQ data signal line 112 (see, e.g., triangles on the DQ and DQS lines of FIGS. 3A and 3B which point towards the memory controller). While various figures of the present application denote read operations by use of DQ and DQS lines which have triangles pointing towards the memory controller, certain embodiments described herein are also compatible with write operations (e.g., as would be denoted by triangles on the DQ or DQS lines pointing away from the memory controller).

For example, in certain embodiments, the circuit 40 comprises a pair of switches 120a, 120b on the DQ data signal lines 102a, 102b as schematically illustrated by FIG. 3A. Each switch 120a, 120b is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ signal line 112, the DQ data signal line 102b to the common DQ signal line 112, or both DQ data signal lines 102a, 102b to the common DQ signal line 112. In certain other embodiments, the circuit 40 comprises a switch 120 electrically coupled to both of the DQ data signal lines 102a, 102b, as schematically illustrated by FIG. 3B. The switch 120 is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ signal line 112,

8

the DQ data signal line 102b to the common DQ signal line 112, or both DQ signal lines 102a, 102b to the common DQ signal line 112. Circuits 40 having other configurations of switches are also compatible with embodiments described herein. While each of the memory devices 30a, 30b has a plurality of DQ data signal lines and a plurality of DQS data strobe signal lines, FIGS. 3A and 3B only illustrate a single DQ data signal line and a single DQS data strobe signal line for each memory device 30a, 30b for simplicity. The configurations schematically illustrated by FIGS. 3A and 3B can be applied to all of the DQ data signal lines and DQS data strobe signal lines of the memory module 10.

In certain embodiments, the circuit 40 selectively isolates the loads of ranks of memory devices 30 from the computer system. As schematically illustrated in FIGS. 4A and 4B, example memory modules 10 compatible with certain embodiments described herein comprise a first number of memory devices 30 arranged in a first number of ranks 32. The memory modules 10 of FIGS. 4A and 4B comprises two ranks 32a, 32b, with each rank 32a, 32b having a corresponding set of DQ data signal lines and a corresponding set of DQS data strobe lines. Other numbers of ranks (e.g., four ranks) of memory devices 30 of the memory module 10 are also compatible with certain embodiments described herein. For simplicity, FIGS. 4A and 4B illustrate only a single DQ data signal line and a single DQS data strobe signal line from each rank 32.

The circuit 40 of FIG. 4A selectively isolates one or more of the DQ data signal lines 102a, 102b of the two ranks 32a, 32b from the computer system. Thus, the circuit 40 selectively allows a DQ data signal to be transmitted from the memory controller 20 of the computer system to the memory devices 30 of one or both of the ranks 32a, 32b via the DQ data signal lines 102a, 102b. In addition, the circuit 40 selectively allows one of a first DQ data signal from the DQ data signal line 102a of the first rank 32a and a second DQ data signal from the DQ data signal line 102b of the second rank 32b to be transmitted to the memory controller 20 via the common DQ data signal line 112. For example, in certain embodiments, the circuit 40 comprises a pair of switches 120a, 120b on the DQ data signal lines 102a, 102b as schematically illustrated by FIG. 4A. Each switch 120a, 120b is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ data signal line 112, the DQ data signal line 102b to the common DQ data signal line 112, or both DQ data signal lines 102a, 102b to the common DQ data signal line 112. In certain other embodiments, the circuit 40 comprises a switch 120 electrically coupled to both of the DQ data signal lines 102a, 102b, as schematically illustrated by FIG. 4B. The switch 120 is selectively actuated to selectively electrically couple the DQ data signal line 102a to the common DQ data signal line 112, the DQ data signal line 102b to the common DQ data signal line 112, or both DQ data signal lines 102a, 102b to the common DQ data signal line 112.

Circuits 40 having other configurations of switches are also compatible with embodiments described herein.

In the example embodiments schematically illustrated by FIGS. 3A, 3B, 4A, and 4B, the circuit 40 comprises a logic element which is integral with and comprises the switches 120 which are coupled to the DQ data signal lines and the DQS data strobe signal lines. In certain such embodiments, each switch 120 comprises a data path multiplexer/demultiplexer. In certain other embodiments, the circuit 40 comprises a logic element 122 which is a separate component operatively coupled to the switches 120, as schematically illustrated by FIGS. 5A-5D. The one or more switches 120 are operatively coupled to the logic element 122 to receive con-

US 8,081,536 B1

9

trol signals from the logic element 122 and to selectively electrically couple one or more data signal lines to a common data signal line. Example switches compatible with embodiments described herein include, but are not limited to field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex. Example logic elements 122 compatible with certain embodiments described herein include, but are not limited to, programmable-logic devices (PLD), application-specific integrated circuits (ASIC), field-programmable gate arrays (FPGA), custom-designed semiconductor devices, and complex programmable-logic devices (CPLD). Example logic elements 122 are available from Lattice Semiconductor Corporation of Hillsboro, Oreg., Altera Corporation of San Jose, Calif., and Xilinx Incorporated of San Jose, Calif.

In certain embodiments, the load isolation provided by the circuit 40 advantageously allows the memory module 10 to present a reduced load (e.g., electrical load, such as capacitive load, inductive load, or impedance load) to the computer system by selectively switching between the two ranks of memory devices 30 to which it is coupled. This feature is used in certain embodiments in which the load of the memory module 10 may otherwise limit the number of ranks or the number of memory devices per memory module. In certain embodiments, the memory module 10 operates as having a data path rank buffer which advantageously isolates the ranks of memory devices 30 of the memory module 10 from one another, from the ranks on other memory modules, and from the computer system. This data path rank buffer of certain embodiments advantageously provides DQ-DQS paths for each rank or sets of ranks of memory devices which are separate from one another, or which are separate from the memory controller of the computer system. In certain embodiments, the load isolation advantageously diminishes

10

the effects of capacitive loading, jitter and other sources of noise. In certain embodiments, the load isolation advantageously simplifies various other aspects of operation of the memory module 10, including but not limited to, setup-and-hold time, clock skew, package skew, and process, temperature, voltage, and transmission line variations.

For certain memory module applications that utilize multiple ranks of memory, increased load on the memory bus can degrade speed performance. In certain embodiments described herein, selectively isolating the loads of the ranks of memory devices 30 advantageously decreases the load on the computer system, thereby allowing the computer system (e.g., server) to run faster with improved signal integrity. In certain embodiments, load isolation advantageously provides system memory with reduced electrical loading, thereby improving the electrical topology to the memory controller 20. In certain such embodiments, the speed and the memory density of the computer system are advantageously increased without sacrificing one for the other.

20 In certain embodiments, load isolation advantageously increases the size of the memory array supported by the memory controller 20 of the computer system. The larger memory array has an increased number of memory devices 30 and ranks of memory devices 30 of the memory module 10, with a corresponding increased number of chip selects. Certain embodiments described herein advantageously provide more system memory using fewer chip selects, thereby avoiding the chip select limitation of the memory controller.

An exemplary section of Verilog code corresponding to logic compatible with a circuit 40 which provides load isolation is listed below in Example 1. The exemplary code of Example 1 corresponds to a circuit 40 comprising six FET switches for providing load isolation to DQ and DQS lines.

Example 1

```

//===== declarations
reg rasN_R, casN_R, weN_R;
wire actv_cmd_R, pch_cmd_R, pch_all_cmd_R, ap_xfr_cmd_R_R;
wire xfr_cmd_R, mrs_cmd_R, rd_cmd_R;
// ----- DDR 2 FET
reg brs0N_R; // registered chip sel
reg brs1N_R; // registered chip sel
reg brs2N_R; // registered chip sel
reg brs3N_R; // registered chip sel
wire sel;
wire sel_01;
wire sel_23;
wire rd_R1;
wire wr_cmd_R_wr_R1;
reg rd_R2_rd_R3_rd_R4_rd_R5;
reg wr_R2_wr_R3_wr_R4_wr_R5;
reg enfet1,enfet2,enfet3,enfet4,enfet5,enfet6;
reg wr_01_R1,wr_23_R1;
reg wr_01_R2,wr_01_R3,wr_01_R4;
reg wr_23_R2,wr_23_R3,wr_23_R4;
wire rodt0_a,rodt0_b;
//===== logic
always @ (posedge clk_in)
begin
    brs0N_R <= brs0_in_N; // cs0
    brs1N_R <= brs1_in_N; // cs1
    brs2N_R <= brs2_in_N; // cs2
    brs3N_R <= brs3_in_N; // cs3
    rasN_R <= brras_in_N;
    casN_R <= brcas_in_N;
    weN_R <= bw_in_N;
end
assign sel = ~brs0N_R | ~brs1N_R | ~brs2N_R | ~brs3N_R ;
assign sel_01 = ~brs0N_R | ~brs1N_R ;
assign sel_23 = ~brs2N_R | ~brs3N_R ;

```

US 8,081,536 B1

11

12

-continued

```

assign activ_cmd_R = !rasN_R & casN_R & weN_R; // activate cmd
assign pch_cmd_R = !rasN_R & casN_R & !weN_R; // pchg cmd
assign xfr_cmd_R = rasN_R & !casN_R; // xfr cmd
assign mrs_cmd = !rasN_R & !casN_R & tweN_R; // md reg set cmd
assign rd_cmd_R = rasN_R & !casN_R & weN_R; // read cmd
assign wr_cmd_R = rasN_R & !casN_R & !weN_R; // write cmd
//-----
assign rd_R1 = sel & rd_cmd_R; // rd cmd cyc 1
assign wr_R1 = sel & wr_cmd_R; // wr cmd cyc 1
//-----
always @(posedge clk_in)
begin
    rd_R2 <= rd_R1;
    rd_R3 <= rd_R2;
    rd_R4 <= rd_R3;
    rd_R5 <= rd_R4;
    rd0_o_R6 <= rd0_o_R5;
    wr_R2 <= wr_R1;
    wr_R3 <= wr_R2;
    wr_R4 <= wr_R3;
    wr_R5 <= wr_R4;
end
//-----
assign wr_01_R1 = sel_01 & wr_cmd_R; // wr cmd cyc 1 for cs 2 & cs3
assign wr_23_R1 = sel_23 & wr_cmd_R; // wr cmd cyc 1 for cs 2 & cs3
always @(posedge clk_in)
begin
    wr_01_R2 <= wr_01_R1;
    wr_01_R3 <= wr_01_R2;
    wr_01_R4 <= wr_01_R3;
    wr_23_R2 <= wr_23_R1;
    wr_23_R3 <= wr_23_R2;
    wr_23_R4 <= wr_23_R3;
end
assign n_rdt0_ab = (rdt0) // odt cmd from sys
| (wr_23_R1) // wr 1st cyc to other mks (assume single dimm per channel)
| (wr_23_R2) // wr 2nd cyc to other mks (assume single dimm per channel)
| (wr_23_R3) // wr 3rd cyc to other mks (assume single dimm per channel)
;
assign rdt1_ab = (rdt1) // odt cmd from sys
| (wr_01_R1) // wr 1st cyc to other mks (assume single dimm per channel)
| (wr_01_R2) // wr 2nd cyc to other mks (assume single dimm per channel)
| (wr_01_R3) // wr 3rd cyc to other mks (assume single dimm per channel)
;
//-----
always @(posedge clk_in)
begin
    if (
        | (rd_R2) // pre-am rd
        | (rd_R3) // 1st cyc of rd brst (c13)
        | (rd_R4) // 2nd cyc of rd brst (c13)
        | (wr_R1) // pre-am wr
        | (wr_R2) // wr brst 1st cyc
        | (wr_R3) // wr brst 2nd cyc
    ) begin
        enfet1 <= 1'b1; // enable fet
        enfet2 <= 1'b1; // enable fet
        enfet3 <= 1'b1; // enable fet
        enfet4 <= 1'b1; // enable fet
        enfet5 <= 1'b1; // enable fet
        enfet6 <= 1'b1; // enable fet
    end
    else begin
        enfet1 <= 1'b0; // disable fet
        enfet2 <= 1'b0; // disable fet
        enfet3 <= 1'b0; // disable fet
        enfet4 <= 1'b0; // disable fet
        enfet5 <= 1'b0; // disable fet
        enfet6 <= 1'b0; // disable fet
    end
end

```

Back-to-Back Adjacent Read Commands

Due to their source synchronous nature, DDR SDRAM (e.g., DDR1, DDR2, DDR3) memory devices operate with a data transfer protocol which surrounds each burst of data strobes with a pre-amble time interval and a post-amble time

65

interval. The pre-amble time interval provides a timing window for the receiving memory device to enable its data capture circuitry when a known valid level is present on the strobe signal to avoid false triggers of the memory device's capture circuit. The post-amble time interval provides extra time after

US 8,081,536 B1

13

the last strobe for this data capture to facilitate good signal integrity. In certain embodiments, when the computer system accesses two consecutive bursts of data from the same memory device, termed herein as a “back-to-back adjacent read,” the post-amble time interval of the first read command and the pre-amble time interval of the second read command are skipped by design protocol to increase read efficiency. FIG. 6A shows an exemplary timing diagram of this “gapless” read burst for a back-to-back adjacent read condition from one memory device.

In certain embodiments, when the second read command accesses data from a different memory device than does the first read command, there is at least one time interval (e.g., clock cycle) inserted between the data strobes of the two memory devices. This inserted time interval allows both read data bursts to occur without the post-amble time interval of the first read data burst colliding or otherwise interfering with the pre-amble time interval of the second read data burst. In certain embodiments, the memory controller of the computer system inserts an extra clock cycle between successive read commands issued to different memory devices, as shown in the exemplary timing diagram of FIG. 6B for successive read accesses from different memory devices.

In typical computer systems, the memory controller is informed of the memory boundaries between the ranks of memory of the memory module prior to issuing read commands to the memory module. Such memory controllers can insert wait time intervals or clock cycles to avoid collisions or interference between back-to-back adjacent read commands which cross memory device boundaries, which are referred to herein as “BBARX.”

In certain embodiments described herein in which the number of ranks 32 of the memory module 10 is doubled or quadrupled, the circuit 40 generates a set of output address and command signals so that the selection decoding is transparent to the computer system. However, in certain such embodiments, there are memory device boundaries of which the computer system is unaware, so there are occasions in which BBARX occurs without the cognizance of the memory controller 20 of the computer system. As shown in FIG. 7, the last data strobe of memory device “a” collides with the preamble time interval of the data strobe of memory device “b,” resulting in a “collision window.”

FIGS. 8A-8D schematically illustrate circuit diagrams of example memory modules 10 comprising a circuit 40 which multiplexes the DQS data strobe signal lines 104a, 104b of two ranks 32a, 32b from one another in accordance with certain embodiments described herein. While the DQS data strobe signal lines 104a, 104b of FIGS. 8A-8D correspond to two ranks 32a, 32b of memory devices 30, in certain other embodiments, the circuit 40 multiplexes the DQS data strobe signal lines 104a, 104b corresponding to two individual memory devices 30a, 30b.

FIG. 8A schematically illustrates a circuit diagram of an exemplary memory module 10 comprising a circuit 40 in accordance with certain embodiments described herein. In certain embodiments, BBARX collisions are avoided by a mechanism which electrically isolates the DQS data strobe signal lines 104a, 104b from one another during the transition from the first read data burst of one rank 32a of memory devices 30 to the second read data burst of another rank 32b of memory devices 30.

In certain embodiments, as schematically illustrated by FIG. 8A, the circuit 40 comprises a first switch 130a electrically coupled to a first DQS data strobe signal line 104a of a first rank 32a of memory devices 30 and a second switch 130b electrically coupled to a second DQS data strobe signal line

14

104b of a second rank 32b of memory devices 30. In certain embodiments, the time for switching the first switch 130a and the second switch 130b is between the two read data bursts (e.g., after the last DQS data strobe of the read data burst of the first rank 32a and before the first DQS data strobe of the read data burst of the second rank 32b). During the read data burst for the first rank 32a, the first switch 130a is enabled. After the last DQS data strobe of the first rank 32a and before the first DQS data strobe of the second rank 32b, the first switch 130a is disabled and the second switch 130b is enabled.

As shown in FIG. 8A, each of the ranks 32a, 32b otherwise involved in a BBARX collision have their DQS data strobe signal lines 104a, 104b selectively electrically coupled to the common DQS line 114 through the circuit 40. The circuit 40 of certain embodiments multiplexes the DQS data strobe signal lines 104a, 104b of the two ranks 32a, 32b of memory devices 30 from one another to avoid a BBARX collision.

In certain embodiments, as schematically illustrated by FIG. 8B, the circuit 40 comprises a switch 130 which multiplexes the DQS data strobe signal lines 104a, 104b from one another. For example, the circuit 40 receives a DQS data strobe signal from the common DQS data strobe signal line 114 and selectively transmits the DQS data strobe signal to the first DQS data strobe signal line 104a, to the second DQS data strobe signal line 104b, or to both DQS data strobe signal lines 104a, 104b. As another example, the circuit 40 receives a first DQS data strobe signal from the first rank 32a of memory devices 30 and a second DQS data strobe signal from a second rank 32b of memory devices 30 and selectively switches one of the first and second DQS data strobe signals to the common DQS data strobe signal line 114.

In certain embodiments, the circuit 40 also provides the load isolation described above in reference to FIGS. 1-5. For example, as schematically illustrated by FIG. 8C, the circuit 40 comprises both the switch 120 for the DQ data signal lines 102a, 102b and the switch 130 for the DQS data strobe signal lines 104a, 104b. While in certain embodiments, the switches 130 are integral with a logic element of the circuit 40, in certain other embodiments, the switches 130 are separate components which are operatively coupled to a logic element 122 of the circuit 40, as schematically illustrated by FIG. 8D. In certain such embodiments, the control and timing of the switch 130 is performed by the circuit 40 which is resident on the memory module 10. Example switches 130 compatible with embodiments described herein include, but are not limited to field-effect transistor (FET) switches, such as the SN74AUC1G66 single bilateral analog switch available from Texas Instruments, Inc. of Dallas, Tex., and multiplexers, such as the SN74AUC2G53 2:1 analog multiplexer/demultiplexer available from Texas Instruments, Inc. of Dallas, Tex.

The circuit 40 of certain embodiments controls the isolation of the DQS data strobe signal lines 104a, 104b by monitoring commands received by the memory module 10 from the computer system and producing “windows” of operation whereby the appropriate switches 130 are activated or deactivated to enable and disable the DQS data strobe signal lines 104a, 104b to mitigate BBARX collisions. In certain other embodiments, the circuit 40 monitors the commands received by the memory module 10 from the computer system and selectively activates or deactivates the switches 120 to enable and disable the DQ data signal lines 102a, 102b to reduce the load of the memory module 10 on the computer system. In still other embodiments, the circuit 40 performs both of these functions together.

Command Signal Translation

Most high-density memory modules are currently built with 512-Megabit (“512-Mb”) memory devices wherein each

US 8,081,536 B1

15

memory device has a 64M×8-bit configuration. For example, a 1-Gigabyte (“1-GB”) memory module with error checking capabilities can be fabricated using eighteen such 512-Mb memory devices. Alternatively, it can be economically advantageous to fabricate a 1-GB memory module using lower-density memory devices and doubling the number of memory devices used to produce the desired word width. For example, by fabricating a 1-GB memory module using thirty-six 256-Mb memory devices with 64M×4-bit configuration, the cost of the resulting 1-GB memory module can be reduced since the unit cost of each 256-Mb memory device is typically lower than one-half the unit cost of each 512-Mb memory device. The cost savings can be significant, even though twice as many 256-Mb memory devices are used in place of the 512-Mb memory devices. For example, by using pairs of 512-Mb memory devices rather than single 1-Gb memory devices, certain embodiments described herein reduce the cost of the memory module by a factor of up to approximately five.

Market pricing factors for DRAM devices are such that higher-density DRAM devices (e.g., 1-Gb DRAM devices) are much more than twice the price of lower-density DRAM devices (e.g., 512-Mb DRAM devices). In other words, the price per bit ratio of the higher-density DRAM devices is greater than that of the lower-density DRAM devices. This pricing difference often lasts for months or even years after the introduction of the higher-density DRAM devices, until volume production factors reduce the costs of the newer higher-density DRAM devices. Thus, when the cost of a higher-density DRAM device is more than the cost of two lower-density DRAM devices, there is an economic incentive for utilizing pairs of the lower-density DRAM devices to replace individual higher-density DRAM devices.

FIG. 9A schematically illustrates an exemplary memory module 10 compatible with certain embodiments described herein. The memory module 10 is connectable to a memory controller 20 of a computer system (not shown). The memory module 10 comprises a printed circuit board 210 and a plurality of memory devices 30 coupled to the printed circuit board 210. The plurality of memory devices 30 has a first number of memory devices 30. The memory module 10 further comprises a circuit 40 coupled to the printed circuit board 210. The circuit 40 receives a set of input address and command signals from the computer system. The set of input address and command signals correspond to a second number of memory devices 30 smaller than the first number of memory devices 30. The circuit 40 generates a set of output address and command signals in response to the set of input address and command signals. The set of output address and command signals corresponds to the first number of memory devices 30.

In certain embodiments, as schematically illustrated in FIG. 9A, the memory module 10 further comprises a phase-lock loop device 220 coupled to the printed circuit board 210 and a register 230 coupled to the printed circuit board 210. In certain embodiments, the phase-lock loop device 220 and the register 230 are each mounted on the printed circuit board 210. In response to signals received from the computer system, the phase-lock loop device 220 transmits clock signals to the plurality of memory devices 30, the circuit 40, and the register 230. The register 230 receives and buffers a plurality of command signals and address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select signals), and transmits corresponding signals to the appropriate memory devices 30. In certain embodiments, the register 230 comprises a plurality of register devices. While the phase-lock

16

loop device 220, the register 230, and the circuit 40 are described herein in certain embodiments as being separate components, in certain other embodiments, two or more of the phase-lock loop device 220, the register 230, and the circuit 40 are portions of a single component. Persons skilled in the art are able to select a phase-lock loop device 220 and a register 230 compatible with embodiments described herein.

In certain embodiments, the memory module 10 further comprises electrical components which are electrically coupled to one another and are surface-mounted or embedded on the printed circuit board 210. These electrical components can include, but are not limited to, electrical conduits, resistors, capacitors, inductors, and transistors. In certain embodiments, at least some of these electrical components are discrete, while in other certain embodiments, at least some of these electrical components are constituents of one or more integrated circuits.

In certain embodiments, the printed circuit board 210 is mountable in a module slot of the computer system. The printed circuit board 210 of certain such embodiments has a plurality of edge connections electrically coupled to corresponding contacts of the module slot and to the various components of the memory module 10, thereby providing electrical connections between the computer system and the components of the memory module 10.

In certain embodiments, the plurality of memory devices 30 are arranged in a first number of ranks 32. For example, in certain embodiments, the memory devices 30 are arranged in four ranks 32a, 32b, 32c, 32d, as schematically illustrated by FIG. 9A. In certain other embodiments, the memory devices 30 are arranged in two ranks 32a, 32b, as schematically illustrated by FIG. 9B. Other numbers of ranks 32 of the memory devices 30 are also compatible with embodiments described herein.

As schematically illustrated by FIGS. 9A and 9B, in certain embodiments, the circuit 40 receives a set of input command signals (e.g., refresh, precharge) and address signals (e.g., bank address signals, row address signals, column address signals, gated column address strobe signals, chip-select signals) from the memory controller 20 of the computer system. In response to the set of input address and command signals, the circuit 40 generates a set of output address and command signals.

In certain embodiments, the set of output address and command signals corresponds to a first number of ranks in which the plurality of memory devices 30 of the memory module 10 are arranged, and the set of input address and command signals corresponds to a second number of ranks per memory module for which the computer system is configured. The second number of ranks in certain embodiments is smaller than the first number of ranks. For example, in the exemplary embodiment as schematically illustrated by FIG. 9A, the first number of ranks is four while the second number of ranks is two. In the exemplary embodiment of FIG. 9B, the first number of ranks is two while the second number of ranks is one. Thus, in certain embodiments, even though the memory module 10 actually has the first number of ranks of memory devices 30, the memory module 10 simulates a virtual memory module by operating as having the second number of ranks of memory devices 30. In certain embodiments, the memory module 10 simulates a virtual memory module when the number of memory devices 30 of the memory module 10 is larger than the number of memory devices 30 per memory module for which the computer system is configured to utilize. In certain embodiments, the circuit 40 comprises logic (e.g., address decoding logic, command decoding logic)

US 8,081,536 B1

17

which translates between a system memory domain of the computer system and a physical memory domain of the memory module 10.

In certain embodiments, the computer system is configured for a number of ranks per memory module which is smaller than the number of ranks in which the memory devices 30 of the memory module 10 are arranged. In certain such embodiments, the computer system is configured for two ranks of memory per memory module (providing two chip-select signals CS₀, CS₁) and the plurality of memory modules 30 of the memory module 10 are arranged in four ranks, as schematically illustrated by FIG. 9A. In certain other such embodiments, the computer system is configured for one rank of memory per memory module (providing one chip-select signal CS₀) and the plurality of memory modules 30 of the memory module 10 are arranged in two ranks, as schematically illustrated by FIG. 9B.

In the exemplary embodiment schematically illustrated by FIG. 9A, the memory module 10 has four ranks of memory devices 30 and the computer system is configured for two ranks of memory devices per memory module. The memory module 10 receives row/column address signals or signal bits (A₀-A_{n+1}), bank address signals (BA₀-BA_m), chip-select signals (CS₀ and CS₁), and command signals (e.g., refresh, precharge, etc.) from the computer system. The A₀-A_n row/column address signals are received by the register 230, which buffers these address signals and sends these address signals to the appropriate ranks of memory devices 30. The circuit 40 receives the two chip-select signals (CS₀, CS₁) and one row/column address signal (A_{n+1}) from the computer system. Both the circuit 40 and the register 230 receive the bank address signals (BA₀-BA_m) and at least one command signal (e.g., refresh, precharge, etc.) from the computer system.

Logic Tables

Table 1 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices 30 using chip-select signals.

TABLE 1

State	CS ₀	CS ₁	A _{n+1}	Command	CS _{0A}	CS _{0B}	CS _{1A}	CS _{1B}
1	0	1	0	Active	0	1	1	1
2	0	1	1	Active	1	0	1	1
3	0	1	x	Active	0	0	1	1
4	1	0	0	Active	1	1	0	1
5	1	0	1	Active	1	1	1	0
6	1	0	x	Active	1	1	0	0
7	1	1	x	x	1	1	1	1

Note:

1. CS₀, CS₁, CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are active low signals.

2. A_{n+1} is an active high signal.

3. 'x' is a Don't Care condition.

4. Command involves a number of command signals that define operations such as refresh, precharge, and other operations.

In Logic State 1: CS₀ is active low, A_{n+1} is non-active, and Command is active. CS_{0A} is pulled low, thereby selecting Rank 0.

In Logic State 2: CS₀ is active low, A_{n+1} is active, and Command is active. CS_{0B} is pulled low, thereby selecting Rank 1.

In Logic State 3: CS₀ is active low, A_{n+1} is Don't Care, and Command is active high. CS_{0A} and CS_{0B} are pulled low, thereby selecting Ranks 0 and 1.

In Logic State 4: CS₁ is active low, A_{n+1} is non-active, and Command is active. CS_{1A} is pulled low, thereby selecting Rank 2.

18

In Logic State 5: CS₁ is active low, A_{n+1} is active, and Command is active. CS_{1B} is pulled low, thereby selecting Rank 3.

In Logic State 6: CS₁ is active low, A_{n+1} is Don't Care, and Command is active. CS_{1A} and CS_{1B} are pulled low, thereby selecting Ranks 2 and 3.

In Logic State 7: CS₀ and CS₁ are pulled non-active high, which deselects all ranks, i.e., CS_{0A}, CS_{0B}, CS_{1A}, and CS_{1B} are pulled high.

The "Command" column of Table 1 represents the various commands that a memory device (e.g., a DRAM device) can execute, examples of which include, but are not limited to, activation, read, write, precharge, and refresh. In certain embodiments, the command signal is passed through to the selected rank only (e.g., state 4 of Table 1). In such embodiments, the command signal (e.g., read) is sent to only one memory device or the other memory device so that data is supplied from one memory device at a time. In other embodiments, the command signal is passed through to both associated ranks (e.g., state 6 of Table 1). In such embodiments, the command signal (e.g., refresh) is sent to both memory devices to ensure that the memory content of the memory devices remains valid over time. Certain embodiments utilize a logic table such as that of Table 1 to simulate a single memory device from two memory devices by selecting two ranks concurrently.

Table 2 provides a logic table compatible with certain embodiments described herein for the selection among ranks of memory devices 30 using gated CAS signals.

TABLE 2

	CS*	RAS*	CAS*	WE*	Density Bit	A ₁₀	Command	CAS0*	CAS1*
35	1	x	x	x	x	x	NOP	x	x
	0	1	1	1	x	x	NOP	1	1
	0	0	1	1	0	x	ACTIVATE	1	1
	0	0	1	1	1	x	ACTIVATE	1	1
	0	1	0	1	0	x	READ	0	1
40	0	1	0	1	1	x	READ	1	0
	0	1	0	0	0	x	WRITE	0	1
	0	1	0	0	1	x	WRITE	1	0
	0	0	1	0	0	0	PRE-CHARGE	1	1
	0	0	1	0	1	0	PRE-CHARGE	1	1
45	0	0	1	0	x	1	PRE-CHARGE	1	1
	0	0	0	0	x	x	MODE	0	0
	0	0	0	1	x	x	REG SET		
	0	0	0	1	x	x	REFRESH	0	0

In certain embodiments in which the density bit is a row address bit, for read/write commands, the density bit is the value latched during the activate command for the selected bank.

Serial-Presence-Detect Device

Memory modules typically include a serial-presence detect (SPD) device 240 (e.g., an electrically-erasable-programmable read-only memory or EEPROM device) comprising data which characterize various attributes of the memory module, including but not limited to, the number of row addresses, the number of column addresses, the data width of the memory devices, the number of ranks, the memory density per rank, the number of memory devices, and the memory density per memory device. The SPD device 240 communicates this data to the basic input/output system (BIOS) of the computer system so that the computer system is informed of the memory capacity and the memory configuration available

US 8,081,536 B1

19

for use and can configure the memory controller properly for maximum reliability and performance.

For example, for a commercially-available 512-MB (64M \times 8-byte) memory module utilizing eight 512-Mb memory devices each with a 64M \times 8-bit configuration, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3:	Defines the number of row address bits in the DRAM device in the memory module [13 for the 512-Mb memory device].
Byte 4:	Defines the number of column address bits in the DRAM device in the memory module [11 for the 512-Mb memory device].
Byte 13:	Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 512-Mb (64 M \times 8-bit) memory device].
Byte 14:	Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 512-Mb (64 M \times 8-bit) memory device].
Byte 17:	Defines the number of banks internal to the DRAM device used in the memory module [4 for the 512-Mb memory device].

In a further example, for a commercially-available 1-GB (128M \times 8-byte) memory module utilizing eight 1-Gb memory devices each with a 128M \times 8-bit configuration, as described above, the SPD device contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3:	Defines the number of row address bits in the DRAM device in the memory module [14 for the 1-Gb memory device].
Byte 4:	Defines the number of column address bits in the DRAM device in the memory module [11 for the 1-Gb memory device].
Byte 13:	Defines the bit width of the primary DRAM device used in the memory module [8 bits for the 1-Gb (128 M \times 8-bit) memory device].
Byte 14:	Defines the bit width of the error checking DRAM device used in the memory module [8 bits for the 1-Gb (128 M \times 8-bit) memory device].
Byte 17:	Defines the number of banks internal to the DRAM device used in the memory module [4 for the 1-Gb memory device].

In certain embodiments, the SPD device 240 comprises data which characterize the memory module 10 as having fewer ranks of memory devices than the memory module 10 actually has, with each of these ranks having more memory density. For example, for a memory module 10 compatible with certain embodiments described herein having two ranks of memory devices 30, the SPD device 240 comprises data which characterizes the memory module 10 as having one rank of memory devices with twice the memory density per rank. Similarly, for a memory module 10 compatible with certain embodiments described herein having four ranks of memory devices 30, the SPD device 240 comprises data which characterizes the memory module 10 as having two ranks of memory devices with twice the memory density per rank. In addition, in certain embodiments, the SPD device 240 comprises data which characterize the memory module 10 as having fewer memory devices than the memory module 10 actually has, with each of these memory devices having more memory density per memory device. For example, for a memory module 10 compatible with certain embodiments described herein, the SPD device 240 comprises data which characterizes the memory module 10 as having one-half the number of memory devices that the memory module 10 actually has, with each of these memory devices having twice the memory density per memory device. Thus, in certain embodiments,

20

the SPD device 240 informs the computer system of the larger memory array by reporting a memory device density that is a multiple of the memory devices 30 resident on the memory module 10. Certain embodiments described herein advantageously do not require system level changes to hardware (e.g., the motherboard of the computer system) or to software (e.g., the BIOS of the computer system).

FIG. 9C schematically illustrates an exemplary memory module 10 in accordance with certain embodiments described herein. The memory module 10 comprises a pair of substantially identical memory devices 31, 33. Each memory device 31, 33 has a first bit width, a first number of banks of memory locations, a first number of rows of memory locations, and a first number of columns of memory locations. The memory module 10 further comprises an SPD device 240 comprising data that characterizes the pair of memory devices 31, 33. The data characterize the pair of memory devices 31, 33 as a virtual memory device having a second bit width equal to twice the first bit width, a second number of banks of memory locations equal to the first number of banks, a second number of rows of memory locations equal to the first number of rows, and a second number of columns of memory locations equal to the first number of columns.

In certain such embodiments, the SPD device 240 of the memory module 10 is programmed to describe the combined pair of lower-density memory devices 31, 33 as one virtual or pseudo-higher-density memory device. In an exemplary embodiment, two 512-Mb memory devices, each with a 128M \times 4-bit configuration, are used to simulate one 1-Gb memory device having a 128M \times 8-bit configuration. The SPD device 240 of the memory module 10 is programmed to describe the pair of 512-Mb memory devices as one virtual or pseudo-1-Gb memory device.

For example, to fabricate a 1-GB (128M \times 8-byte) memory module, sixteen 512-Mb (128M \times 4-bit) memory devices can be used. The sixteen 512-Mb (128M \times 4-bit) memory devices are combined in eight pairs, with each pair serving as a virtual or pseudo-1-Gb (128M \times 8-bit) memory device. In certain such embodiments, the SPD device 240 contains the following SPD data (in appropriate bit fields of these bytes):

Byte 3:	13 row address bits.
Byte 4:	12 column address bits.
Byte 13:	8 bits wide for the primary virtual 1-Gb (128 M \times 8-bit) memory device.
Byte 14:	8 bits wide for the error checking virtual 1-Gb (128 M \times 8-bit) memory device.
Byte 17:	4 banks.

In this exemplary embodiment, bytes 3, 4, and 17 are programmed to have the same values as they would have for a 512-MB (128M \times 4-byte) memory module utilizing 512-Mb (128M \times 4-bit) memory devices. However, bytes 13 and 14 of the SPD data are programmed to be equal to 8, corresponding to the bit width of the virtual or pseudo-higher-density 1-Gb (128M \times 8-bit) memory device, for a total capacity of 1-GB. Thus, the SPD data does not describe the actual-lower-density memory devices, but instead describes the virtual or pseudo-higher-density memory devices. The BIOS accesses the SPD data and recognizes the memory module as having 4 banks of memory locations arranged in 2^{13} rows and 2^{12} columns, with each memory location having a width of 8 bits rather than 4 bits.

In certain embodiments, when such a memory module 10 is inserted in a computer system, the computer system's memory controller then provides to the memory module 10 a

US 8,081,536 B1

21

set of input address and command signals which correspond to the number of ranks or the number of memory devices reported by the SPD device **240**. For example, placing a two-rank memory module **10** compatible with certain embodiments described herein in a computer system compatible with one-rank memory modules, the SPD device **240** reports to the computer system that the memory module **10** only has one rank. The circuit **40** then receives a set of input address and command signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output address and command signals corresponding to two ranks to the appropriate memory devices **30** of the memory module **10**.

Similarly, when a two-rank memory module **10** compatible with certain embodiments described herein is placed in a computer system compatible with either one- or two-rank memory modules, the SPD device **240** reports to the computer system that the memory module **10** only has one rank. The circuit **40** then receives a set of input address and command signals corresponding to a single rank from the computer system's memory controller, and generates and transmits a set of output address and command signals corresponding to two ranks to the appropriate memory devices **30** of the memory module **10**.

Furthermore, a four-rank memory module **10** compatible with certain embodiments described herein simulates a two-rank memory module whether the memory module **10** is inserted in a computer system compatible with two-rank memory modules or with two- or four-rank memory modules. Thus, by placing a four-rank memory module **10** compatible with certain embodiments described herein in a module slot that is four-rank-ready, the computer system provides four chip-select signals, but the memory module **10** only uses two of the chip-select signals.

In certain embodiments, the circuit **40** comprises the SPD device **240** which reports the CAS latency (CL) to the memory controller of the computer system. The SPD device **240** of certain embodiments reports a CL which has one more cycle than does the actual operational CL of the memory array. In certain embodiments, data transfers between the memory controller and the memory module are registered for one additional clock cycle by the circuit **40**. The additional clock cycle of certain embodiments is added to the transfer time budget with an incremental overall CAS latency. This extra cycle of time in certain embodiments advantageously provides sufficient time budget to add a buffer which electrically isolates the ranks of memory devices **30** from the memory controller **20**. The buffer of certain embodiments comprises combinatorial logic, registers, and logic pipelines. In certain embodiments, the buffer adds a one-clock cycle time delay, which is equivalent to a registered DIMM, to accomplish the address decoding. The one-cycle time delay of certain such embodiments provides sufficient time for read and write data transfers to provide the functions of the data path multiplexer/demultiplexer. Thus, for example, a DDR2 400-MHz memory system in accordance with embodiments described herein has an overall CAS latency of four, and uses memory devices with a CAS latency of three. In still other embodiments, the SPD device **240** does not utilize this extra cycle of time.

Memory Density Multiplication

In certain embodiments, two memory devices having a memory density are used to simulate a single memory device having twice the memory density, and an additional address signal bit is used to access the additional memory. Similarly, in certain embodiments, two ranks of memory devices having

22

a memory density are used to simulate a single rank of memory devices having twice the memory density, and an additional address signal bit is used to access the additional memory. As used herein, such simulations of memory devices or ranks of memory devices are termed as "memory density multiplication," and the term "density transition bit" is used to refer to the additional address signal bit which is used to access the additional memory by selecting which rank of memory devices is enabled for a read or write transfer operation.

For example, for computer systems which are normally limited to using memory modules which have a single rank of 128M \times 4-bit memory devices, certain embodiments described herein enable the computer system to utilize memory modules which have double the memory (e.g., two ranks of 128M \times 4-bit memory devices). The circuit **40** of certain such embodiments provides the logic (e.g., command and address decoding logic) to double the number of chip selects, and the SPD device **240** reports a memory device density of 256M \times 4-bit to the computer system.

In certain embodiments utilizing memory density multiplication embodiments, the memory module **10** can have various types of memory devices **30** (e.g., DDR1, DDR2, DDR3, and beyond). The circuit **40** of certain such embodiments utilizes implied translation logic equations having variations depending on whether the density transition bit is a row, column, or internal bank address bit. In addition, the translation logic equations of certain embodiments vary depending on the type of memory module **10** (e.g., UDIMM, RDIMM, FBDIMM, etc.). Furthermore, in certain embodiments, the translation logic equations vary depending on whether the implementation multiplies memory devices per rank or multiplies the number of ranks per memory module.

Table 3A provides the numbers of rows and columns for DDR-1 memory devices, as specified by JEDEC standard JESD79D, "Double Data Rate (DDR) SDRAM Specification," published February 2004, and incorporated in its entirety by reference herein.

TABLE 3A

	128-Mb	256-Mb	512-Mb	1-Gb
Number of banks	4	4	4	4
Number of row address bits	12	13	13	14
Number of column address bits for "x 4" configuration	11	11	12	12
Number of column address bits for "x 8" configuration	10	10	11	11
Number of column address bits for "x 16" configuration	9	9	10	10

As described by Table 3A, 512-Mb (128M \times 4-bit) DRAM devices have 2^{13} rows and 2^{12} columns of memory locations, while 1-Gb (128M \times 8-bit) DRAM devices have 2^{14} rows and 2^{11} columns of memory locations. Because of the differences in the number of rows and the number of columns for the two types of memory devices, complex address translation procedures and structures would typically be needed to fabricate a 1-GB (128M \times 8-byte) memory module using sixteen 512-Mb (128M \times 4-bit) DRAM devices.

US 8,081,536 B1

23

Table 3B shows the device configurations as a function of memory density for DDR2 memory devices.

TABLE 3B

Number of Rows	Number of Columns	Number of Internal Banks	Page Size (x4s or x8s)
256 Mb	13	11	4
512 Mb	14	11	4
1 Gb	14	11	8
2 Gb	15	11	8
4 Gb	16	11	8

Table 4 lists the corresponding density transition bit for the density transitions between the DDR2 memory densities of Table 3B.

TABLE 4

Density Transition	Density Transition Bit
256 Mb to 512 Mb	A ₁₃
512 Mb to 1 Gb	BA ₂
1 Gb to 2 Gb	A ₁₄
2 Gb to 4Gb	A ₁₅

Other certain embodiments described herein utilize a transition bit to provide a transition from pairs of physical 4-Gb memory devices to simulated 8-Gb memory devices.

In an example embodiment, the memory module comprises one or more pairs of 256-Mb memory devices, with each pair simulating a single 512-Mb memory device. The simulated 512-Mb memory device has four internal banks while each of the two 256-Mb memory devices has four internal banks, for a total of eight internal banks for the pair of 256-Mb memory devices. In certain embodiments, the additional row address bit is translated by the circuit 40 to the rank selection between each of the two 256-Mb memory devices of the pair. Although there are eight total internal banks in the rank-converted memory array, the computer system is only aware of four internal banks. When the memory controller activates a row for a selected bank, the circuit 40 activates the same row for the same bank, but it does so for the selected rank according to the logic state of the additional row address bit A₁₃.

In another example embodiment, the memory module comprises one or more pairs of 512-Mb memory devices, with each pair simulating a single 1-Gb memory device. The simulated 1-Gb memory device has eight internal banks while each of the two 512-Mb memory devices has four internal banks, for a total of eight internal banks for the pair of 512-Mb

24

memory devices. In certain embodiments, the mapped BA₂ (bank 2) bit is used to select between the two ranks of 512-Mb memory devices to preserve the internal bank geometry expected by the memory controller of the computer system. The state of the BA₂ bit selects the upper or lower set of four banks, as well as the upper and lower 512-Mb rank.

In another example embodiment, the memory module comprises one or more pairs of 1-Gb memory devices, with each pair simulating a single 2-Gb memory device. Each of the two 1-Gb memory devices has eight internal banks for a total of sixteen internal banks, while the simulated 2-Gb memory device has eight internal banks. In certain embodiments, the additional row address bit translates to the rank selection between the two 1-Gb memory devices. Although there are sixteen total internal banks per pair of 1-Gb memory devices in the rank-converted memory array, the memory controller of the computer system is only aware of eight internal banks. When the memory controller activates a row of a selected bank, the circuit 40 activates the same row for the same bank, but it does so for the selected rank according to the logic state of the additional row address bit A₁₄.

The circuit 40 of certain embodiments provides substantially all of the translation logic used for the decoding (e.g., command and address decoding). In certain such embodiments, there is a fully transparent operational conversion from the “system memory” density domain of the computer system to the “physical memory” density domain of the memory module 10. In certain embodiments, the logic translation equations are programmed in the circuit 40 by hardware, while in certain other embodiments, the logic translation equations are programmed in the circuit 40 by software. Examples 1 and 2 provide exemplary sections of Verilog code compatible with certain embodiments described herein. As described more fully below, the code of Examples 1 and 2 includes logic to reduce potential problems due to “back-to-back adjacent read commands which cross memory device boundaries or “BBARX.” Persons skilled in the art are able to provide additional logic translation equations compatible with embodiments described herein.

An exemplary section of Verilog code compatible with memory density multiplication from 512 Mb to 1 Gb using DDR2 memory devices with the BA₂ density transition bit is listed below in Example 2. The exemplary code of Example 2 corresponds to a circuit 40 which receives one chip-select signal from the computer system and which generates two chip-select signals.

Example 2

```

always @(posedge clk_in)
begin
    rs0N_R <= rs0M_N; // cs0
    rasN_R <= ras_in_N;
    casN_R <= cas_in_N;
    weN_R <= we_in_N;
end
// Gated Chip Selects
assign
    pcs0a_1 = (~rs0_in_N & ~ras_in_N & ~cas_in_N) // ref,md reg set
    | (~rs0_in_N & ras_in_N & cas_in_N) // ref exit, pwr dn
    | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in) // pchg all
    | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & ~a10_in & ~ba2_in) // pchg single bnk
    | (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N & ~ba2_in) // activate
    | (~rs0_in_N & ras_in_N & ~cas_in_N & ~we_in_N & ~ba2_in) // xfr
    ;

```

US 8,081,536 B1

25

26

-continued

```

assign          pcs0b_1 = (~rs0_in_N & ~ras_in_N & ~cas_in_N)           // ref,md reg set
               | (~rs0_in_N & ras_in_N & cas_in_N)           // ref exit, pwr dn
               | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & a10_in)      // pchg all
               | (~rs0_in_N & ~ras_in_N & cas_in_N & ~we_in_N & ~a10_in & ba2_in)    // pchg single bnk
               | (~rs0_in_N & ~ras_in_N & cas_in_N & we_in_N & ba2_in)           // activate
               | (~rs0_in_N & ras_in_N & ~cas_in_N)           & ba2_in) // xfr
               ;
//-----
always @(posedge clk_in)
begin
  a4_r <= a4_in ;
  a5_r <= a5_in ;
  a6_r <= a6_in ;
  a10_r <= a10_in ;
  ba0_r <= ba0_in ;
  ba1_r <= ba1_in ;
  ba2_r <= ba2_in ;
  q_mrs_cmd_cyc1 <= q_mrs_cmd ;
end
///////////////////////////////
// determine the cas latency
///////////////////////////////
assign q_mrs_cmd_r = (!rasN_R & !casN_R & !weN_R)
  & !rs0N_R
  & (!ba0_r & !ba1_r)
  ;
always @(posedge clk_in) // lmr
  if (~reset_N) // lmr
    c13 <= 1'b1 ;
  else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
  c13 <= (~a6_r & a5_r & a4_r) ;
end
always @(posedge clk_in)
  if (~reset_N) // reset
    c12 <= 1'b0 ;
  else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
  c12 <= (~a6_r & a5_r & ~a4_r) ;
end
always @(posedge clk_in)
  if (~reset_N) c15 <= 1'b0 ;
  else if (q_mrs_cmd_cyc1) // load mode reg cmd
begin
  c15 <= (a6_r & ~a5_r & a4_r) ;
end
assign          pre_cyc2_enfet = (wr_cmd_cyc1 & acs_cyc1 & c13)           // wr brst c13 preamble
               ;
assign          pre_cyc3_enfet = (rd_cmd_cyc2 & c13)           // rd brst c13 preamble
               | (wr_cmd_cyc2 & c13)           // wr brst c13 1st pair
               | (wr_cmd_cyc2 & c14)           // wr brst c14 preamble
               ;
assign          pre_cyc4_enfet = (wr_cmd_cyc3 & c13)           // wr brst c13 2nd pair
               | (wr_cmd_cyc3 & c14)           // wr brst c14 1st pair
               | (rd_cmd_cyc3 & c13)           // rd brst c13 1st pair
               | (rd_cmd_cyc3 & c14)           // rd brst c14 preamble
               ;
assign          pre_cyc5_enfet = (rd_cmd_cyc4 & c13)           // rd brst c13 2nd pair
               | (wr_cmd_cyc4 & c14)           // wr brst c14 2nd pair
               | (rd_cmd_cyc4 & c14)           // rd brst c14 1st pair
               ;
// dq
assign          pre_dq_cyc = pre_cyc2_enfet
               | pre_cyc3_enfet
               | pre_cyc4_enfet
               | pre_cyc5_enfet
               ;
assign          pre_dq_ncyc = enfet_cyc2
               | enfet_cyc3
               | enfet_cyc4
               | enfet_cyc5
               ;

```

US 8,081,536 B1

27

28

-continued

```
// dqs
assign      pre_dqsa_cyc = (pre_cyc2_enfet & ~ba2_r)
            | (pre_cyc3_enfet & ~ba2_cyc2)
            | (pre_cyc4_enfet & ~ba2_cyc3)
            | (pre_cyc5_enfet & ~ba2_cyc4)
            ;
assign      pre_dqsb_cyc = (pre_cyc2_enfet & ba2_r)
            | (pre_cyc3_enfet & ba2_cyc2)
            | (pre_cyc4_enfet & ba2_cyc3)
            | (pre_cyc5_enfet & ba2_cyc4)
            ;
assign      pre_dqsa_ncyc = (enfet_cyc2 & ~ba2_cyc2)
            | (enfet_cyc3 & ~ba2_cyc3)
            | (enfet_cyc4 & ~ba2_cyc4)
            | (enfet_cyc5 & ~ba2_cyc5)
            ;
assign      pre_dqsb_ncyc = (enfet_cyc2 & ba2_cyc2)
            | (enfet_cyc3 & ba2_cyc3)
            | (enfet_cyc4 & ba2_cyc4)
            | (enfet_cyc5 & ba2_cyc5)
            ;
always @(posedge clk_in)
begin
    acs_cyc2 <= acs_cyc1;           // cs active
    ba2_cyc2 <= ba2_r;
    ba2_cyc3 <= ba2_cyc2;
    ba2_cyc4 <= ba2_cyc3;
    ba2_cyc5 <= ba2_cyc4;
    rd_cmd_cyc2 <= rd_cmd_cyc1 & acs_cyc1;
    rd_cmd_cyc3 <= rd_cmd_cyc2 ;
    rd_cmd_cyc4 <= rd_cmd_cyc3 ;
    rd_cmd_cyc5 <= rd_cmd_cyc4 ;
    rd_cmd_cyc6 <= rd_cmd_cyc5 ;
    rd_cmd_cyc7 <= rd_cmd_cyc6 ;
    wr_cmd_cyc2 <= wr_cmd_cyc1 & acs_cyc1;
    wr_crnd_cyc3 <= wr_cmd_cyc2 ;
    wr_cmd_cyc4 <= wr_cmd_cyc3 ;
    wr_cmd_cyc5 <= wr_cmd_cyc4 ;
end
always @ (negedge clk_in)
begin
    dq_ncyc <= dq_cyc;
    dqs_ncyc_a <= dqs_cyc_a;
    dqs_ncyc_b <= dqs_cyc_b;
end
// DQ FET enables
assign      enq_fet1 = dq_cyc | dq_ncyc      ;
assign      enq_fet2 = dq_cyc | dq_ncyc ;
assign      enq_fet3 = dq_cyc | dq_ncyc ;
assign      enq_fet4 = dq_cyc | dq_ncyc ;
assign      enq_fet5 = dq_cyc | dq_ncyc ;
// DQS FET enables
assign      ens_fet1a = dqs_cyc_a | dqs_ncyc_a ;
assign      ens_fet2a = dqs_cyc_a | dqs_ncyc_a ;
assign      ens_fet3a = dqs_cyc_a | dqs_ncyc_a ;
assign      ens_fet1b = dqs_cyc_b | dqs_ncyc_b ;
assign      ens_fet2b = dqs_cyc_b | dqs_ncyc_b ;
assign      ens_fet3b = dqs_cyc_b | dqs_ncyc_b ;
```

Another exemplary section of Verilog code compatible with memory density multiplication from 256 Mb to 512 Mb using DDR2 memory devices and gated CAS signals with the row A₁₃ density transition bit is listed below in Example 3. The exemplary code of Example 3 corresponds to a circuit 40

which receives one gated CAS signal from the computer system and which generates two gated CAS signals.

Example 3

```
// latched a13 flags cs0, banks 0-3
always @(posedge clk_in)
if (actv_cmd_R & ~ra0N_R & ~bnk1_R & ~bnk0_R) // activate
begin
    1_a13_00 <= a13_r;
end
```

US 8,081,536 B1

29

30

-continued

```

always @@(posedge clk_in)
if (actv_cmd_R & ~rs0N_R & ~bnk1_R & bnk0_R)      // activate
begin
    1_a13_01 <= a13_r;
end
always @@(posedge clk_in)
if (actv_cmd_R & ~r0N_R & bnlk1_R & ~bnk0_R)      // activate
begin
    1_a13_10 <= a13_r;
end
always @@(posedge clk_in)
if (actv_cmd_R & ~rs0N_R & bnlk1_R & bnlk0_R)      // activate
begin
    1_a13_11 <= a13_r;
end
// gated cas
assign cas_i = ~(casN_R);
assign cas0_o = (~rasN_R & cas_i)
| (rasN_R & ~1_a13_00 & ~bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & ~1_a13_01 & ~bnk1_R & bnlk0_R & cas_i)
| (rasN_R & ~1_a13_10 & bnlk1_R & ~bnk0_R & cas_i)
| (rasN_R & ~1_a13_11 & bnlk1_R & bnlk0_R & cas_i)
;
assign cas1_o = (~rasN_R & cas_i)
| (rasN_R & 1_a13_00      & ~bnk1_R & ~bnk0_R & cas_i)
| (rasN_R & 1_a13_01      & ~bnk1_R & bnlk0_R & cas_i)
| (rasN_R & 1_a13_10      & bnlk1_R & ~bnk0_R & cas_i)
| (rasN_R & 1_a13_11      & bnlk1_R & bnlk0_R & cas_i)
;
assign pcas_0_N = ~cas0_o;
assign pcas_1_N = ~cas1_o;
assign rd0_o_R1 = rasN_R & cas0_o & weN_R & ~rs0N_R; // mk0 rd cmd cyc
assign rd1_o_R1 = rasN_R & cas1_o & weN_R & ~rs0N_R; // mk1 rd cmd cyc
assign wr0_o_R1 = rasN_R & cas0_o & ~weN_R & ~rs0N_R; // mk0 wr cmd cyc
assign wr1_o_R1 = rasN_R & cas1_o & ~weN_R & ~rs0N_R; // mk1 wr cmd cyc
always @@(posedge clk_in)
begin
    rd0_o_R2 <= rd0_o_R1 ;
    rd0_o_R3 <= rd0_o_R2;
    rd0_o_R4 <= rd0_o_R3;
    rd0_o_R5 <= rd0_o_R4;
    rd1_o_R2 <= rd1_o_R1 ;
    rd1_o_R3 <= rd1_o_R2;
    rd1_o_R4 <= rd1_o_R3;
    rd1_o_R5 <= rd1_o_R4;
    wr0_o_R2 <= wr0_o_R1 ;
    wr0_o_R3 <= wr0_o_R2;
    wr0_o_R4 <= wr0_o_R3;
    wr1_o_R2 <= wr1_o_R1 ;
    wr1_o_R3 <= wr1_o_R2;
    wr1_o_R4 <= wr1_o_R3;
end
always @@(posedge clk_in)
begin
    if(
        (rd0_o_R2 & ~rd1_o_R4)                                // pre-am rd if no ped on rnk 1
        | rd0_o_R3                                         // 1st cyc of rd brst
        | rd0_o_R4                                         // 2nd cyc of rd brst
        | (rd0_o_R5 & ~rd1_o_R2 & ~rd1_o_R3)           // post-rd cyc if no ped on rnk 1
        | (wr0_o_R1)                                       // pre-am wr
        | wr0_o_R2 | wr0_o_R3                            // wr brst 1st & 2nd cyc
        | (wr0_o_R4)                                       // post-wr cyc (chgef9)
        | wr1_o_R1 | wr1_o_R2 | wr1_o_R3 | wr1_o_R4 // rank 1 (chgef9)
    )
        en_fet_a <= 1'b1;                                     // enable fet
    else
        en_fet_a <= 1'b0;                                     // disable fet
    end
always @@(posedge clk_in)
begin
    if(
        (rd1_o_R2 & ~rd0_o_R4)
        | rd1_o_R3
        | rd1_o_R4
        | (rd1_o_R5 & ~rd0_o_R2 & ~rd0_o_R3)           // (chgef8)
        | (wr1_o_R1)
        | wr1_o_R2 | wr1_o_R3
    )

```

US 8,081,536 B1

31

-continued

```

| (wr1_o_R4)
| wr0_o_R1 | wr0_o_R2 | wr0_o_R3 | wr0_o_R4
|
)           en_fet_b <= 1'b1;
else         en_fet_b <= 1'b0;
end

```

32

In certain embodiments, the chipset memory controller of the computer system uses the inherent behavioral characteristics of the memory devices (e.g., DDR2 memory devices) to optimize throughput of the memory system. For example, for each internal bank in the memory array, a row (e.g., 1 KB page) is advantageously held activated for an extended period of time. The memory controller, by anticipating a high number of memory accesses or hits to a particular region of memory, can exercise this feature to advantageously eliminate time-consuming pre-charge cycles. In certain such embodiments in which two half-density memory devices are transparently substituted for a single full-density memory device (as reported by the SPD device 240 to the memory controller), the memory devices advantageously support the “open row” feature.

FIG. 10A schematically illustrates an exemplary memory module 10 which doubles the rank density in accordance with certain embodiments described herein. The memory module 10 has a first memory capacity. The memory module 10 comprises a plurality of substantially identical memory devices 30 configured as a first rank 32a and a second rank 32b. In certain embodiments, the memory devices 30 of the first rank 32a are configured in pairs, and the memory devices 30 of the second rank 32b are also configured in pairs. In certain embodiments, the memory devices 30 of the first rank 32a are configured with their respective DQS pins tied together and the memory devices 30 of the second rank 32b are configured with their respective DQS pins tied together, as described more fully below. The memory module 10 further comprises a circuit 40 which receives a first set of address and command signals from a memory controller (not shown) of the computer system. The first set of address and command signals is compatible with a second memory capacity substantially equal to one-half of the first memory capacity. The circuit 40 translates the first set of address and command signals into a second set of address and command signals which is compatible with the first memory capacity of the memory module 10 and which is transmitted to the first rank 32a and the second rank 32b.

The first rank 32a of FIG. 10A has 18 memory devices 30 and the second rank 32b of FIG. 10A has 18 memory devices 30. Other numbers of memory devices 30 in each of the ranks 32a, 32b are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. 10A, the memory module 10 has a width of 8 bytes (or 64 bits) and each of the memory devices 30 of FIG. 10A has a bit width of 4 bits. The 4-bit-wide (“×4”) memory devices 30 of FIG. 10A have one-half the width, but twice the depth of 8-bit-wide (“×8”) memory devices. Thus, each pair of “×4” memory devices 30 has the same density as a single “×8” memory device, and pairs of “×4” memory devices 30 can be used instead of individual “×8” memory devices to provide the memory density of the memory module 10. For example, a pair of 512-Mb 128M×4-bit memory devices has the same memory density as a 1-Gb 128M×8-bit memory device.

10 For two “×4” memory devices 30 to work in tandem to mimic a “×8” memory device, the relative DQS pins of the two memory devices 30 in certain embodiments are advantageously tied together, as described more fully below. In addition, to access the memory density of a high-density memory module 10 comprising pairs of “×4” memory devices 30, an additional address line is used. While a high-density memory module comprising individual “×8” memory devices with the next-higher density would also utilize an additional address line, the additional address lines are different in the two memory module configurations.

For example, a 1-Gb 128M×8-bit DDR-1 DRAM memory device uses row addresses A₁₃-A₀ and column addresses A₁₁ and A₉-A₀. A pair of 512-Mb 128M×4-bit DDR-1 DRAM 25 memory devices uses row addresses A₁₂-A₀ and column addresses A₁₂, A₁₁, and A₉-A₀. In certain embodiments, a memory controller of a computer system utilizing a 1-GB 128M×8 memory module 10 comprising pairs of the 512-Mb 128M×4 memory devices 30 supplies the address and command signals including the extra row address (A₁₃) to the memory module 10. The circuit 40 receives the address and command signals from the memory controller and converts the extra row address (A₁₃) into an extra column address (A₁₂).

35 FIG. 10B schematically illustrates an exemplary circuit 40 compatible with embodiments described herein. The circuit 40 is used for a memory module 10 comprising pairs of “×4” memory devices 30 which mimic individual “×8” memory devices. In certain embodiments, each pair has the respective DQS pins of the memory devices 30 tied together. In certain embodiments, as schematically illustrated by FIG. 10B, the circuit 40 comprises a programmable-logic device (PLD) 42, a first multiplexer 44 electrically coupled to the first rank 32a of memory devices 30, and a second multiplexer 46 electrically 45 coupled to the second rank 32b of memory devices 30. In certain embodiments, the PLD 42 and the first and second multiplexers 44, 46 are discrete elements, while in other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can select an appropriate PLD 42, first multiplexer 44, and second multiplexer 46 50 in accordance with embodiments described herein.

In the exemplary circuit 40 of FIG. 10B, during a row access procedure (CAS is high), the first multiplexer 44 passes the A₁₂ address through to the first rank 32, the second multiplexer 46 passes the A₁₂ address through to the second rank 34, and the PLD 42 saves or latches the A₁₃ address from the memory controller. In certain embodiments, a copy of the A₁₃ address is saved by the PLD 42 for each of the internal banks (e.g., 4 internal banks) per memory device 30. During 55 a subsequent column access procedure (CAS is low), the first multiplexer 44 passes the previously-saved A₁₃ address through to the first rank 32a as the A₁₂ address and the second multiplexer 46 passes the previously-saved A₁₃ address through to the second rank 32b as the A₁₂ address. The first rank 32a and the second rank 32b thus interpret the previously-saved A₁₃ row address as the current A₁₂ column address. In this way, in certain embodiments, the circuit 40 60

US 8,081,536 B1

33

translates the extra row address into an extra column address in accordance with certain embodiments described herein.

Thus, by allowing two lower-density memory devices to be used rather than one higher-density memory device, certain embodiments described herein provide the advantage of using lower-cost, lower-density memory devices to build “next-generation” higher-density memory modules. Certain embodiments advantageously allow the use of lower-cost readily-available 512-Mb DDR-2 SDRAM devices to replace more expensive 1-Gb DDR-2 SDRAM devices. Certain embodiments advantageously reduce the total cost of the resultant memory module.

FIG. 11A schematically illustrates an exemplary memory module 10 which doubles number of ranks in accordance with certain embodiments described herein. The memory module 10 has a first plurality of memory locations with a first memory density. The memory module 10 comprises a plurality of substantially identical memory devices 30 configured as a first rank 32a, a second rank 32b, a third rank 32c, and a fourth rank 32d. The memory module 10 further comprises a circuit 40 which receives a first set of address and command signals from a memory controller (not shown). The first set of address and command signals is compatible with a second plurality of memory locations having a second memory density. The second memory density is substantially equal to one-half of the first memory density. The circuit 40 translates the first set of address and command signals into a second set of address and command signals which is compatible with the first plurality of memory locations of the memory module 10 and which is transmitted to the first rank 32a, the second rank 32b, the third rank 32c, and the fourth rank 32d.

Each rank 32a, 32b, 32c, 32d of FIG. 11A has 9 memory devices 30. Other numbers of memory devices 30 in each of the ranks 32a, 32b, 32c, 32d are also compatible with embodiments described herein.

In the embodiment schematically illustrated by FIG. 11A, the memory module 10 has a width of 8 bytes (or 64 bits) and each of the memory devices 30 of FIG. 11A has a bit width of 8 bits. Because the memory module 10 has twice the number of 8-bit-wide (“ $\times 8$ ”) memory devices 30 as does a standard 8-byte-wide memory module, the memory module 10 has twice the density as does a standard 8-byte-wide memory module. For example, a 1-GB 128M \times 8-byte memory module with 36 512-Mb 128M \times 8-bit memory devices (arranged in four ranks) has twice the memory density as a 512-Mb 128M \times 8-byte memory module with 18 512-Mb 128M \times 8-bit memory devices (arranged in two ranks).

To access the additional memory density of the high-density memory module 10, the two chip-select signals (CS₀, CS₁) are used with other address and command signals to gate a set of four gated CAS signals. For example, to access the additional ranks of four-rank 1-GB 128M \times 8-byte DDR-1 DRAM memory module, the CS₀ and CS₁ signals along with the other address and command signals are used to gate the CAS signal appropriately, as schematically illustrated by FIG. 11A. FIG. 11B schematically illustrates an exemplary circuit 40 compatible with embodiments described herein. In certain embodiments, the circuit 40 comprises a programmable-logic device (PLD) 42 and four “OR” logic elements 52, 54, 56, 58 electrically coupled to corresponding ranks 32a, 32b, 32c, 32d of memory devices 30.

In certain embodiments, the PLD 42 comprises an ASIC, an FPGA, a custom-designed semiconductor device, or a CPLD. In certain embodiments, the PLD 42 and the four “OR” logic elements 52, 54, 56, 58 are discrete elements, while in other certain embodiments, they are integrated within a single integrated circuit. Persons skilled in the art can

34

select an appropriate PLD 42 and appropriate “OR” logic elements 52, 54, 56, 58 in accordance with embodiments described herein.

In the embodiment schematically illustrated by FIG. 11B, the PLD 42 transmits each of the four “enabled CAS” (ENCAS_{0a}, ENCAS_{0b}, ENCAS_{1a}, ENCAS_{1b}) signals to a corresponding one of the “OR” logic elements 52, 54, 56, 58. The CAS signal is also transmitted to each of the four “OR” logic elements 52, 54, 56, 58. The CAS signal and the “enabled CAS” signals are “low” true signals. By selectively activating each of the four “enabled CAS” signals which are inputted into the four “OR” logic elements 52, 54, 56, 58, the PLD 42 is able to select which of the four ranks 32a, 32b, 32c, 32d is active.

In certain embodiments, the PLD 42 uses sequential and combinatorial logic procedures to produce the gated CAS signals which are each transmitted to a corresponding one of the four ranks 32a, 32b, 32c, 32d. In certain other embodiments, the PLD 42 instead uses sequential and combinatorial logic procedures to produce four gated chip-select signals (e.g., CS_{0a}, CS_{0b}, CS_{1a}, and CS_{1b}) which are each transmitted to a corresponding one of the four ranks 32a, 32b, 32c, 32d.

Tied Data Strobe Signal Pins

For proper operation, the computer system advantageously recognizes a 1-GB memory module comprising 256-Mb memory devices with 64M \times 4-bit configuration as a 1-GB memory module having 512-Mb memory devices with 64M \times 8-bit configuration (e.g., as a 1-GB memory module with 128M \times 8-byte configuration). This advantageous result is desirably achieved in certain embodiments by electrically connecting together two output signal pins (e.g., DQS or data strobe pins) of the two 256-Mb memory devices such that both output signal pins are concurrently active when the two memory devices are concurrently enabled. The DQS or data strobe is a bi-directional signal that is used during both read cycles and write cycles to validate or latch data. As used herein, the terms “tying together” or “tied together” refer to a configuration in which corresponding pins (e.g., DQS pins) of two memory devices are electrically connected together and are concurrently active when the two memory devices are concurrently enabled (e.g., by a common chip-select or CS signal). Such a configuration is different from standard memory module configurations in which the output signal pins (e.g., DQS pins) of two memory devices are electrically coupled to the same source, but these pins are not concurrently active since the memory devices are not concurrently enabled. However, a general guideline of memory module design warns against tying together two output signal pins in this way.

FIGS. 12 and 13 schematically illustrate a problem which may arise from tying together two output signal pins. FIG. 12 schematically illustrates an exemplary memory module 305 in which a first DQS pin 312 of a first memory device 310 is electrically connected to a second DQS pin 322 of a second memory device 320. The two DQS pins 312, 322 are both electrically connected to a memory controller 330.

FIG. 13 is an exemplary timing diagram of the voltages applied to the two DQS pins 312, 322 due to non-simultaneous switching. As illustrated by FIG. 13, at time t₁, both the first DQS pin 312 and the second DQS pin 322 are high, so no current flows between them. Similarly, at time t₄, both the first DQS pin 312 and the second DQS pin 322 are low, so no current flows between them. However, for times between approximately t₂ and approximately t₃, the first DQS pin 312 is low while the second DQS pin 322 is high. Under such conditions, a current will flow between the two DQS pins 312,

US 8,081,536 B1

35

322. This condition in which one DQS pin is low while the other DQS pin is high can occur for fractions of a second (e.g., 0.8 nanoseconds) during the dynamic random-access memory (DRAM) read cycle. During such conditions, the current flowing between the two DQS pins 312, 322 can be substantial, resulting in heating of the memory devices 310, 320, and contributing to the degradation of reliability and eventual failure of these memory devices.

A second problem may also arise from tying together two output signal pins. FIG. 14 schematically illustrates another exemplary memory module 305 in which a first DQS pin 312 of a first memory device 310 is electrically connected to a second DQS pin 322 of a second memory device 320. The two DQS pins 312, 322 of FIG. 14 are both electrically connected to a memory controller (not shown). The DQ (data input/output) pin 314 of the first memory device 310 and the corresponding DQ pin 324 of the second memory device 320 are each electrically connected to the memory controller by the DQ bus (not shown). Typically, each memory device 310, 320 will have a plurality of DQ pins (e.g., eight DQ pins per memory device), but for simplicity, FIG. 14 only shows one DQ pin for each memory device 310, 320.

Each of the memory devices 310, 320 of FIG. 14 utilizes a respective on-die termination or “ODT” circuit 332, 334 which has termination resistors (e.g., 75 ohms) internal to the memory devices 310, 320 to provide signal termination. Each memory device 310, 320 has a corresponding ODT signal pin 362, 364 which is electrically connected to the memory controller via an ODT bus 340. The ODT signal pin 362 of the first memory device 310 receives a signal from the ODT bus 340 and provides the signal to the ODT circuit 332 of the first memory device 310. The ODT circuit 332 responds to the signal by selectively enabling or disabling the internal termination resistors 352, 356 of the first memory device 310. This behavior is shown schematically in FIG. 14 by the switches 342, 344 which are either closed (dash-dot line) or opened (solid line). The ODT signal pin 364 of the second memory device 320 receives a signal from the ODT bus 340 and provides the signal to the ODT circuit 334 of the second memory device 320. The ODT circuit 334 responds to the signal by selectively enabling or disabling the internal termination resistors 354, 358 of the second memory device 320. This behavior is shown schematically in FIG. 14 by the switches 346, 348 which are either closed (dash-dot line) or opened (solid line). The switches 342, 344, 346, 348 of FIG. 14 are schematic representations of the operation of the ODT circuits 332, 334, and do not signify that the ODT circuits 332, 334 necessarily include mechanical switches.

Examples of memory devices 310, 320 which include such ODT circuits 332, 334 include, but are not limited to, DDR2 memory devices. Such memory devices are configured to selectively enable or disable the termination of the memory device in this way in response to signals applied to the ODT signal pin of the memory device. For example, when the ODT signal pin 362 of the first memory device 310 is pulled high, the termination resistors 352, 356 of the first memory device 310 are enabled. When the ODT signal pin 362 of the first memory device 310 is pulled low (e.g., grounded), the termination resistors 352, 356 of the first memory device 310 are disabled. By selectively disabling the termination resistors of an active memory device, while leaving the termination resistors of inactive memory devices enabled, such configurations advantageously preserve signal strength on the active memory device while continuing to eliminate signal reflections at the bus-die interface of the inactive memory devices.

In certain configurations, as schematically illustrated by FIG. 14, the DQS pins 312, 322 of each memory device 310,

36

320 are selectively connected to a voltage VTT through a corresponding termination resistor 352, 354 internal to the corresponding memory device 310, 320. Similarly, in certain configurations, as schematically illustrated by FIG. 14, the DQ pins 314, 324 are selectively connected to a voltage VTT through a corresponding termination resistor 356, 358 internal to the corresponding memory device 310, 320. In certain configurations, rather than being connected to a voltage VTT, the DQ pins 314, 324 and/or the DQS pins 312, 322 are selectively connected to ground through the corresponding termination resistors 352, 354, 356, 358. The resistances of the internal termination resistors 352, 354, 356, 358 are selected to clamp the voltages so as to reduce the signal reflections from the corresponding pins. In the configuration schematically illustrated by FIG. 14, each internal termination resistor 352, 354, 356, 358 has a resistance of approximately 75 ohms.

When connecting the first memory device 310 and the second memory device 320 together to form a double word width, both the first memory device 310 and the second memory device 320 are enabled at the same time (e.g., by a common CS signal). Connecting the first memory device 310 and the second memory device 320 by tying the DQS pins 312, 322 together, as shown in FIG. 14, results in a reduced effective termination resistance for the DQS pins 312, 322. For example, for the exemplary configuration of FIG. 14, the effective termination resistance for the DQS pins 312, 322 is approximately 37.5 ohms, which is one-half the desired ODT resistance (for 75-ohm internal termination resistors) to reduce signal reflections since the internal termination resistors 352, 354 of the two memory devices 310, 320 are connected in parallel. This reduction in the termination resistance can result in signal reflections causing the memory device to malfunction.

FIG. 15 schematically illustrates an exemplary memory module 400 in accordance with certain embodiments described herein. The memory module 400 comprises a first memory device 410 having a first data strobe (DQS) pin 412 and a second memory device 420 having a second data strobe (DQS) pin 422. The memory module 400 further comprises a first resistor 430 electrically coupled to the first DQS pin 412. The memory module 400 further comprises a second resistor 440 electrically coupled to the second DQS pin 422 and to the first resistor 430. The first DQS pin 412 is electrically coupled to the second DQS pin 422 through the first resistor 430 and through the second resistor 440.

In certain embodiments, the memory module 400 is a 1-GB unbuffered Double Data Rate (DDR) Synchronous Dynamic RAM (SDRAM) high-density dual in-line memory module (DIMM). FIGS. 16A and 16B schematically illustrate a first side 462 and a second side 464, respectively, of such a memory module 400 with eighteen 64Mx4-bit, DDR-1 SDRAM FBGA memory devices on each side of a 184-pin glass-epoxy printed circuit board (PCB) 460. In certain embodiments, the memory module 400 further comprises a phase-lock-loop (PLL) clock driver 470, an EEPROM for serial-presence detect (SPD) data 480, and decoupling capacitors (not shown) mounted on the PCB in parallel to suppress switching noise on VDD and VDDQ power supply for DDR-1 SDRAM. By using synchronous design, such memory modules 400 allow precise control of data transfer between the memory module 400 and the system controller. Data transfer can take place on both edges of the DQS signal at various operating frequencies and programming latencies. Therefore, certain such memory modules 400 are suitable for a variety of high-performance system applications.

US 8,081,536 B1

37

In certain embodiments, the memory module 400 comprises a plurality of memory devices configured in pairs, each pair having a first memory device 410 and a second memory device 420. For example, in certain embodiments, a 128Mx-72-bit DDR SDRAM high-density memory module 400 comprises thirty-six 64Mx4-bit DDR-I SDRAM integrated circuits in FBGA packages configured in eighteen pairs. The first memory device 410 of each pair has the first DQS pin 412 electrically coupled to the second DQS pin 422 of the second memory device 420 of the pair. In addition, the first DQS pin 412 and the second DQS pin 422 are concurrently active when the first memory device 410 and the second memory device 420 are concurrently enabled.

In certain embodiments, the first resistor 430 and the second resistor 440 each has a resistance advantageously selected to reduce the current flow between the first DQS pin 412 and the second DQS pin 422 while allowing signals to propagate between the memory controller and the DQS pins 412, 422. In certain embodiments, each of the first resistor 430 and the second resistor 440 has a resistance in a range between approximately 5 ohms and approximately 50 ohms. For example, in certain embodiments, each of the first resistor 430 and the second resistor 440 has a resistance of approximately 22 ohms. Other resistance values for the first resistor 430 and the second resistor 440 are also compatible with embodiments described herein. In certain embodiments, the first resistor 430 comprises a single resistor, while in other embodiments, the first resistor 430 comprises a plurality of resistors electrically coupled together in series and/or in parallel. Similarly, in certain embodiments, the second resistor 440 comprises a single resistor, while in other embodiments, the second resistor 440 comprises a plurality of resistors electrically coupled together in series and/or in parallel.

FIGS. 17A and 17B schematically illustrate an exemplary embodiment of a memory module 400 in which the first resistor 430 and the second resistor 440 are used to reduce the current flow between the first DQS pin 412 and the second DQS pin 422. As schematically illustrated by FIG. 17A, the memory module 400 is part of a computer system 500 having a memory controller 510. The first resistor 430 has a resistance of approximately 22 ohms and the second resistor 440 has a resistance of approximately 22 ohms. The first resistor 430 and the second resistor 440 are electrically coupled in parallel to the memory controller 510 through a signal line 520 having a resistance of approximately 25 ohms. The first resistor 430 and the second resistor 440 are also electrically coupled in parallel to a source of a fixed termination voltage (identified by VTT in FIGS. 17A and 17B) by a signal line 540 having a resistance of approximately 47 ohms. Such an embodiment can advantageously be used to allow two memory devices having lower bit widths (e.g., 4-bit) to behave as a single virtual memory device having a higher bit width (e.g., 8-bit).

FIG. 17B schematically illustrates exemplary current-limiting resistors 430, 440 in conjunction with the impedances of the memory devices 410, 420. During an exemplary portion of a data read operation, the memory controller 510 is in a high-impedance condition, the first memory device 410 drives the first DQS pin 412 high (e.g., 2.7 volts), and the second memory device 420 drives the second DQS pin 422 low (e.g., 0 volts). The amount of time for which this condition occurs is approximated by the time between t₂ and t₃ of FIG. 13, which in certain embodiments is approximately twice the tDQSQ (data strobe edge to output data edge skew time, e.g., approximately 0.8 nanoseconds). At least a portion of this time in certain embodiments is caused by simultaneous switching output (SSO) effects.

38

In certain embodiments, as schematically illustrated by FIG. 17B, the DQS driver of the first memory device 410 has a driver impedance R₁ of approximately 17 ohms, and the DQS driver of the second memory device 420 has a driver impedance R₄ of approximately 17 ohms. Because the upper network of the first memory device 410 and the first resistor 430 (with a resistance R₂ of approximately 22 ohms) is approximately equal to the lower network of the second memory device 420 and the second resistor 440 (with a resistance R₃ of approximately 22 ohms), the voltage at the midpoint is approximately 0.5*(2.7-0)=1.35 volts, which equals VTT, such that the current flow across the 47-ohm resistor of FIG. 17B is approximately zero.

The voltage at the second DQS pin 422 in FIG. 17B is given by V_{DQS2}=2.7*R₄/(R₁+R₂+R₃+R₄)=0.59 volts and the current flowing through the second DQS pin 422 is given by I_{DQS2}=0.59/R₄=34 millamps. The power dissipation in the DQS driver of the second memory device 420 is thus P_{DQS2}=34 mA*0.59 V=20 milliwatts. In contrast, without the first resistor 430 and the second resistor 440, only the 17-ohm impedances of the two memory devices 410, 420 would limit the current flow between the two DQS pins 412, 422, and the power dissipation in the DQS driver of the second memory device 420 would be approximately 107 milliwatts. Therefore, the first resistor 430 and the second resistor 440 of FIGS. 17A and 17B advantageously limit the current flowing between the two memory devices during the time that the DQS pin of one memory device is driven high and the DQS pin of the other memory device is driven low.

In certain embodiments in which there is overshoot or undershoot of the voltages, the amount of current flow can be higher than those expected for nominal voltage values. Therefore, in certain embodiments, the resistances of the first resistor 430 and the second resistor 440 are advantageously selected to account for such overshoot/undershoot of voltages.

For certain such embodiments in which the voltage at the second DQS pin 422 is V_{DQS2}=0.59 volts and the duration of the overdrive condition is approximately 0.8 nanoseconds at maximum, the total surge is approximately 0.59 V*1.2 ns=0.3 V·ns. For comparison, the JEDEC standard for overshoot/undershoot is 2.4 V·ns, so certain embodiments described herein advantageously keep the total surge within predetermined standards (e.g., JEDEC standards).

FIG. 18 schematically illustrates another exemplary memory module 600 compatible with certain embodiments described herein. The memory module 600 comprises a termination bus 605. The memory module 600 further comprises a first memory device 610 having a first data strobe pin 612, a first termination signal pin 614 electrically coupled to the termination bus 605, a first termination circuit 616, and at least one data pin 618. The first termination circuit 616 selectively electrically terminating the first data strobe pin 612 and the first data pin 618 in response to a first signal received by the first termination signal pin 614 from the termination bus 605. The memory module 600 further comprises a second memory device 620 having a second data strobe pin 622 electrically coupled to the first data strobe pin 612, a second termination signal pin 624, a second termination circuit 626, and at least one data pin 628. The second termination signal pin 624 is electrically coupled to a voltage, wherein the second termination circuit 626 is responsive to the voltage by not terminating the second data strobe pin 622 or the second data pin 628. The memory module 600 further comprises at least one termination assembly 630 having a third termination signal pin 634, a third termination circuit 636, and at least one termination pin 638 electrically coupled to the data pin 628 of

US 8,081,536 B1

39

the second memory device 620. The third termination signal pin 634 is electrically coupled to the termination bus 605. The third termination circuit 636 selectively electrically terminates the data pin 628 of the second memory device 620 through the termination pin 638 in response to a second signal received by the third termination signal pin 634 from the termination bus 605.

FIG. 19 schematically illustrates a particular embodiment of the memory module 600 schematically illustrated by FIG. 18. The memory module 600 comprises an on-die termination (ODT) bus 605. The memory module 600 comprises a first memory device 610 having a first data strobe (DQS) pin 612, a first ODT signal pin 614 electrically coupled to the ODT bus 605, a first ODT circuit 616, and at least one data (DQ) pin 618. The first ODT circuit 616 selectively electrically terminates the first DQS pin 612 and the DQ pin 618 of the first memory device 610 in response to an ODT signal received by the first ODT signal pin 614 from the ODT bus 605. This behavior of the first ODT circuit 616 is schematically illustrated in FIG. 14 by the switches 672, 676 which are selectively closed (dash-dot line) or opened (solid line).

The memory module 600 further comprises a second memory device 620 having a second DQS pin 622 electrically coupled to the first DQS pin 612, a second ODT signal pin 624, a second ODT circuit 626, and at least one DQ pin 628. The first DQS pin 612 and the second DQS pin 622 are concurrently active when the first memory device 610 and the second memory device 620 are concurrently enabled. The second ODT signal pin 624 is electrically coupled to a voltage (e.g., ground), wherein the second ODT circuit 626 is responsive to the voltage by not terminating the second DQS pin 622 or the second DQ pin 624. This behavior of the second ODT circuit 626 is schematically illustrated in FIG. 14 by the switches 674, 678 which are opened.

The memory module 600 further comprises at least one termination assembly 630 having a third ODT signal pin 634 electrically coupled to the ODT bus 605, a third ODT circuit 636, and at least one termination pin 638 electrically coupled to the DQ pin 628 of the second memory device 620. The third ODT circuit 636 selectively electrically terminates the DQ pin 628 of the second memory device 620 through the termination pin 638 in response to an ODT signal received by the third ODT signal pin 634 from the ODT bus 605. This behavior of the third ODT circuit 636 is schematically illustrated in FIG. 19 by the switch 680 which is either closed (dash-dot line) or opened (solid line).

In certain embodiments, the termination assembly 630 comprises discrete electrical components which are surface-mounted or embedded on the printed-circuit board of the memory module 600. In certain other embodiments, the termination assembly 630 comprises an integrated circuit mounted on the printed-circuit board of the memory module 600. Persons skilled in the art can provide a termination assembly 630 in accordance with embodiments described herein.

Certain embodiments of the memory module 600 schematically illustrated by FIG. 19 advantageously avoid the problem schematically illustrated by FIG. 12 of electrically connecting the internal termination resistances of the DQS pins of the two memory devices in parallel. As described above in relation to FIG. 14, FIGS. 18 and 19 only show one DQ pin for each memory device for simplicity. Other embodiments have a plurality of DQ pins for each memory device. In certain embodiments, each of the first ODT circuit 616, the second ODT circuit 626, and the third ODT circuit 636 are responsive to a high voltage or signal level by enabling the corresponding termination resistors and are responsive to a

40

low voltage or signal level (e.g., ground) by disabling the corresponding termination resistors. In other embodiments, each of the first ODT circuit 616, the second ODT circuit 626, and the third ODT circuit 636 are responsive to a high voltage or signal level by disabling the corresponding termination resistors and are responsive to a low voltage or signal level (e.g., ground) by enabling the corresponding termination resistors. Furthermore, the switches 672, 674, 676, 678, 680 of FIG. 18 are schematic representations of the enabling and disabling operation of the ODT circuits 616, 626, 636 and do not signify that the ODT circuits 616, 626, 636 necessarily include mechanical switches.

The first ODT signal pin 614 of the first memory device 610 receives an ODT signal from the ODT bus 605. In response to this ODT signal, the first ODT circuit 616 selectively enables or disables the termination resistance for both the first DQS pin 612 and the DQ pin 618 of the first memory device 610. The second ODT signal pin 624 of the second memory device 620 is tied (e.g., directly hard-wired) to the voltage (e.g., ground), thereby disabling the internal termination resistors 654, 658 on the second DQS pin 622 and the second DQ pin 628, respectively, of the second memory device 620 (schematically shown by open switches 674, 678 in FIG. 19). The second DQS pin 622 is electrically coupled to the first DQS pin 612, so the termination resistance for both the first DQS pin 612 and the second DQS pin 622 is provided by the termination resistor 652 internal to the first memory device 510.

The termination resistor 656 of the DQ pin 618 of the first memory device 610 is enabled or disabled by the ODT signal received by the first ODT signal pin 614 of the first memory device 610 from the ODT bus 605. The termination resistance of the DQ pin 628 of the second memory device 620 is enabled or disabled by the ODT signal received by the third ODT signal pin 634 of the termination assembly 630 which is external to the second memory device 620. Thus, in certain embodiments, the first ODT signal pin 614 and the third ODT signal pin 634 receive the same ODT signal from the ODT bus 605, and the termination resistances for both the first memory device 610 and the second memory device 620 are selectively enabled or disabled in response thereto when these memory devices are concurrently enabled. In this way, certain embodiments of the memory module 600 schematically illustrated by FIG. 19 provides external or off-chip termination of the second memory device 620.

Certain embodiments of the memory module 600 schematically illustrated by FIG. 19 advantageously allow the use of two lower-cost readily-available 512-Mb DDR-2 SDRAM devices to provide the capabilities of a more expensive 1-GB DDR-2 SDRAM device. Certain such embodiments advantageously reduce the total cost of the resultant memory module 600.

Certain embodiments described herein advantageously increase the memory capacity or memory density per memory slot or socket on the system board of the computer system. Certain embodiments advantageously allow for higher memory capacity in systems with limited memory slots. Certain embodiments advantageously allow for flexibility in system board design by allowing the memory module 10 to be used with computer systems designed for different numbers of ranks (e.g., either with computer systems designed for two-rank memory modules or with computer systems designed for four-rank memory modules). Certain embodiments advantageously provide lower costs of board designs.

In certain embodiments, the memory density of a memory module is advantageously doubled by providing twice as

US 8,081,536 B1

41

many memory devices as would otherwise be provided. For example, pairs of lower-density memory devices can be substituted for individual higher-density memory devices to reduce costs or to increase performance. As another example, twice the number of memory devices can be used to produce a higher-density memory configuration of the memory module. Each of these examples can be limited by the number of chip select signals which are available from the memory controller or by the size of the memory devices. Certain embodiments described herein advantageously provide a logic mechanism to overcome such limitations.

Various embodiments of the present invention have been described above. Although this invention has been described with reference to these specific embodiments, the descriptions are intended to be illustrative of the invention and are not intended to be limiting. Various modifications and applications may occur to those skilled in the art without departing from the true spirit and scope of the invention.

What is claimed is:

1. A circuit configured to be mounted on a memory module configured to be operationally coupled to a computer system, the memory module having a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of DDR chip-select signals, the circuit including at least one configuration in which the circuit is configured to:

receive a set of signals comprising address signals and a second number of DDR chip-select signals smaller than the first number of DDR chip-select signals;

generate phase-locked clock signals and transmit the phase-locked clock signals to the DDR memory circuits of the first number of ranks;

selectively isolate a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals; and

generate the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of DDR chip-select signals.

2. The circuit of claim 1, wherein at least one address signal of the set of signals is received during a row access procedure and the first number of DDR chip-select signals are generated during a subsequent column address procedure in response at least in part to the at least one address signal received during the row access procedure.

3. The circuit of claim 2, wherein the at least one address signal received during the row access procedure comprises a bank address signal.

4. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to receive an input command signal from the computer system and to generate and transmit an output command signal to the DDR memory circuits of the first number of ranks.

5. The circuit of claim 4, wherein the circuit in the at least one configuration is further configured to respond at least in part to the address signals and the second number of DDR chip-select signals by selecting at least one rank of the first number of ranks and transmitting the output command signal to the DDR memory circuits of the selected at least one rank of the first number of ranks.

6. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to control isolation of data strobe signal lines in response to the set of signals.

42

7. The circuit of claim 1, wherein the circuit has a load and selectively isolating a load of the DDR memory devices of the at least one rank from the computer system comprises presenting the load of the circuit to the computer system.

8. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to selectively electrically couple the DDR memory circuits of the first number of ranks to the computer system and to selectively isolate the DDR memory circuits of the first number of ranks from one another and from the computer system.

9. The circuit of claim 8, wherein the circuit in the at least one configuration is further configured to provide data paths for the DDR memory circuits of each rank of the first number of ranks, wherein the data paths of each rank of the first number of ranks are isolated from one another.

10. The circuit of claim 8, wherein the circuit in the at least one configuration is further configured to provide data paths for the DDR memory circuits of each rank of the first number of ranks, wherein the data paths of each rank of the first number of ranks are isolated from the computer system.

11. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to generate control signals in response to one or more control signals of the set of signals and to transmit the generated control signals to at least some of the DDR memory circuits of the first number of ranks.

12. The circuit of claim 11, wherein the one or more control signals comprises a row address bit and the circuit is further configurable to latch the row address bit.

13. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to transmit a set of output signals to the DDR memory circuits of the first number of ranks.

14. The circuit of claim 13, wherein the circuit in the at least one configuration is further configured to receive command signals from the computer system, to select at least one rank of the first number of ranks, and to transmit the command signals to the DDR memory circuits of the selected at least one rank.

15. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to selectively isolate a data signal line of the DDR memory circuits of the first number of ranks from the computer system.

16. The circuit of claim 1, wherein the memory module has attributes and the circuit in the at least one configuration is further configured to store data accessible to the computer system, wherein the data characterizes the memory module as having attributes that are different from the attributes of the memory module.

17. The circuit of claim 16, wherein the attributes are selected from a group consisting of: a number of row addresses, a number of column addresses, a number of DDR memory circuits, a data width of the DDR memory circuits, a memory density per DDR memory circuit, a number of ranks, and a memory density per rank.

18. The circuit of claim 1, wherein the address signals comprise bank address signals.

19. The circuit of claim 18, wherein the circuit in the at least one configuration is further configured to generate the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the bank address signals, and the second number of DDR chip-select signals.

20. The circuit of claim 18, wherein the set of signals are capable of controlling a memory module having a second number of ranks smaller than the first number of ranks, each

US 8,081,536 B1

43

rank of the second number of ranks comprising DDR memory circuits that are configured to be activated concurrently with one another.

21. The circuit of claim 20, wherein the first number of ranks comprises a first number of DDR memory circuits, and the second number of ranks comprises a second number of DDR memory circuits smaller than the first number of DDR memory circuits.

22. The circuit of claim 1, wherein the circuit in the at least one configuration is further configured to monitor command signals of the set of signals received by the memory module.

23. The circuit of claim 22, wherein the circuit in the at least one configuration is further configured to selectively isolate a load of the DDR memory circuits of the at least one rank of the first number of ranks from the computer system in response at least in part to the command signals.

24. A method of operating a memory module configured to be operationally coupled to a computer system, the memory module having a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of DDR chip-select signals, the method comprising:

receiving a set of signals comprising address signals and a second number of DDR chip-select signals smaller than the first number of DDR chip-select signals;
using the memory module to generate phase-locked clock signals and transmitting the phase-locked clock signals to the DDR memory circuits of the first number of ranks;
selectively isolating a load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals; and

generating the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the address signals, and the second number of DDR chip-select signals.

25. The method of claim 24, further comprising responding at least in part to an address signal received during a row access procedure by generating the first number of DDR chip-select signals during a subsequent column access procedure.

26. The method of claim 24, further comprising receiving an input command signal from the computer system and providing an output command signal to the DDR memory circuits of the first number of ranks.

27. The method of claim 26, further comprising responding at least in part to the address signals and the second number of DDR chip-select signals by selecting at least one rank of the first number of ranks and transmitting the output command signal to the DDR memory circuits of the selected at least one rank of the first number of ranks.

28. The method of claim 24, wherein selectively isolating a load of the DDR memory circuits of the at least one rank from the computer system comprises presenting a smaller load to the computer system.

29. The method of claim 24, further comprising selectively isolating a data signal line of the DDR memory circuits of the first number of ranks from the computer system.

30. The method of claim 24, wherein the memory module has attributes and the method further comprises storing data accessible to the computer system, wherein the data characterizes the memory module as having attributes that are different from the attributes of the memory module.

44

31. The method of claim 30, wherein the attributes are selected from a group consisting of: a number of row addresses, a number of column addresses, a number of DDR memory circuits, a data width of the DDR memory circuits, a memory density per DDR memory circuit, a number of ranks, and a memory density per rank.

32. The method of claim 24, wherein the address signals comprise bank address signals.

33. The method of claim 32, wherein generating the first number of DDR chip-select signals is in response at least in part to the phase-locked clock signals, the bank address signals, and the second number of DDR chip-select signals.

34. The method of claim 24, wherein the set of signals are capable of controlling a memory module with a second number of ranks smaller than the first number of ranks, each rank of the second number of ranks comprising DDR memory circuits that are configured to be activated concurrently with one another.

35. The method of claim 34, wherein the first number of ranks comprises a first number of DDR memory circuits, and the second number of ranks comprises a second number of DDR memory circuits smaller than the first number of DDR memory circuits.

36. The method of claim 24, wherein the method further comprises monitoring command signals of the set of signals received by the memory module.

37. The method of claim 36, wherein selectively isolating a load of the DDR memory circuits of the at least one rank of the first number of ranks from the computer system is performed in response at least in part to the command signals.

38. A circuit configured to be mounted on a memory module configured to be operationally coupled to a computer system, the memory module having a first number of ranks, each rank of the first number of ranks comprising a plurality of double-data-rate (DDR) memory circuits that are configured to be activated concurrently with one another for receiving and transmitting data having a bit width of the rank in response at least in part to a first number of DDR chip-select signals, the circuit including at least one configuration in which the circuit is configured to:

receive a set of signals comprising a command signal, address signals, and a second number of DDR chip-select signals smaller than the first number of DDR chip-select signals;

generate phase-locked clock signals and transmit the phase-locked clock signals to the first number of ranks; respond at least in part to the phase-locked clock signals, address signals, and the second number of DDR chip-select signals by selecting at least one rank of the first number of ranks to receive the command signal and transmitting the command signal to the DDR memory circuits of the selected at least one rank;

selectively isolate a data signal line load of the DDR memory circuits of at least one rank of the first number of ranks from the computer system in response at least in part to the set of signals by presenting a load of the circuit to the computer system; and

receive at least one address signal during a row access procedure and generate the first number of DDR chip-select signals during a subsequent column access procedure in response at least in part to the phase-locked clock signals, the at least one address signal received during the row access procedure, and the second number of DDR chip-select signals.

US 8,081,536 B1

45

39. The circuit of claim **38**, wherein the memory module has attributes and the circuit is further configurable to store data accessible to the computer system, wherein the data characterizes the memory module as having attributes that are different from the attributes of the memory module.

40. The circuit of claim **39**, wherein the attributes are selected from a group consisting of: a number of row addresses, a number of column addresses, a number of DDR memory circuits, a data width of the DDR memory circuits, a memory density per DDR memory circuit, a number of ranks, and a memory density per rank.

46

41. The circuit of claim **38**, wherein the address signals comprise bank address signals.

42. The circuit of claim **41**, wherein the circuit in the at least one configuration is further configured to generate the first number of DDR chip-select signals in response at least in part to the phase-locked clock signals, the bank address signals, and the second number of DDR chip-select signals.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

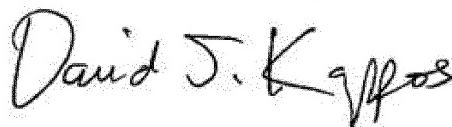
PATENT NO. : 8,081,536 B1
APPLICATION NO. : 13/032470
DATED : December 20, 2011
INVENTOR(S) : Solomon et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

- At Column 1, Line 7 (Approx.), Change “12/995,711,” to --12/955,711,--.
- At Column 1, Line 26, Change “12/995,711,” to --12/955,711,--.
- At Column 2, Line 9-10 (Approx.), Change “128M 8-bit” to --128M×8-bit--.
- At Column 2, Line 31, Change “Elipida” to --Elpida--.
- At Column 11-12, Line 10 (Approx.), Change “assignrd R1” to --assign rd_R 1--.
- At Column 11-12, Line 37, Change “assig n” to --assign--.
- At Column 24, Line 19, Change “is” to --it--.
- At Column 24, Line 35, Change ““back” to --back--.
- At Column 23-24 (Example 2), Line 4 (Approx.), Change “rs0M_N;” to --rs0_in_N;--.
- At Column 25-26, Line 58, Change “(wr cmd_cyc2 & c13)” to --(wr_cmd_cyc2 & c13)--.
- At Column 25-26, Line 59, Change “(wr cmd_cyc2 & c14)” to --(wr_cmd_cyc2 & c14)--.
- At Column 27-28, Line 10, Change “(pre_syc3_enfet & ba2_cyc2)” to --(pre_cyc3_enfet & ba2_cyc2)--.
- At Column 27-28, Line 38, Change “wr_crnd_cyc3” to --wr_cmd_cyc3--.
- At Column 29-30, Line 9, Change “~r0N” to --~rs0N--.

Signed and Sealed this
Fourteenth Day of August, 2012



David J. Kappos
Director of the United States Patent and Trademark Office

CERTIFICATE OF SERVICE

I hereby certify that I electronically filed the foregoing with the Clerk of the Court for the United States Court of Appeals for the Federal Circuit by using the appellate CM/ECF system on August 24, 2016.

I certify that all participants in the case are registered CM/ECF users and that service will be accomplished by the appellate CM/ECF system.

Dated: August 24, 2016

/s/ Brian R. Matsui

CERTIFICATE OF COMPLIANCE WITH RULE 32(a)

This brief complies with the type-volume limitation of Rule 32(a) of the Federal Rules of Appellate Procedure because it contains 12,556 words.

Dated: August 22, 2016

/s/ Brian R. Matsui

dc-845257